arXiv:1602.08048v2 [math.ST] 2 Jan 2017

Denoising Flows on Trees

Sabyasachi Chatterjee* **and John Lafferty*†**

*Department of Statistics
†Department of Computer Science
University of Chicago

January 4, 2017

Abstract: We study the estimation of flows on trees, a structured generalization of isotonic regression. A tree flow is defined recursively as a positive flow value into a node that is partitioned into an outgoing flow to the children nodes, with some amount of the flow possibly leaking outside. We study the behavior of the least squares estimator for flows, and the associated minimax lower bounds. We characterize the risk of the least squares estimator in two regimes. In the first regime the diameter of the tree grows at most logarithmically with the number of nodes. In the second regime, the tree contains many long paths. The results are compared with known risk bounds for isotonic regression.

1. Introduction

We study the problem of denoising tree flows, a graph-structured generalization of isotonic regression. In isotonic (monotonic) regression, a sequence $Y_i = \mu_i + \epsilon_i$, for $i = 1, \ldots, n$, is a noisy observation of a monotonic sequence $\mu_1 \geq \mu_2 \geq \cdots \geq \mu_n$, with ϵ_i independent mean zero random noise. Tree flows are a generalization of monotonic sequences. Each node in a (rooted) tree is labeled with a value that can be thought of as the incoming flow of some fluid. The flow is partitioned and redirected to the children nodes. In a noisy flow, the observation at node i is $Y_i = \mu_i + \epsilon_i$, where μ_i is the true flow and ϵ_i is mean zero random noise.

Figure 1 illustrates a flow on a tree with $n = 9$ nodes. The root node has an incoming flow of 6 units; its left child receives a flow of 3 and its right child receives a flow of 2. The root node thus "leaks" a flow of 1. The node having a flow of 1 is seen to leak a flow of $\frac{1}{6}$. In general, the flow μ_i at a node i having children $\mathcal{C}(i)$ must satisfy $\mu_i \geq 0$ and the flow constraints

$$\mu_i \geq \sum_{j \in \mathcal{C}(i)} \mu_j. \tag{1.1}$$

More explicitly, suppose that the tree T_n has n nodes; the leaves of the tree are the nodes i for which $\mathcal{C}(i)$ is empty. The set of flows $\mathcal{F}(T_n) \subset \mathbb{R}^n$ is the closed convex cone

$$\mathcal{F}(T_n) = \Big\{ \mu \in \mathbb{R}^n \ : \ \mu_i \geq \sum_{j \in \mathcal{C}(i)} \mu_j, \ \text{ for all } i = 1, 2, \ldots, n \Big\}. \tag{1.2}$$

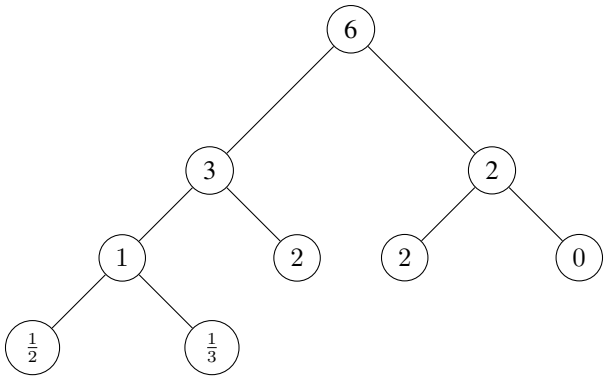

FIG 1. *An example flow. The value μ_j of a node j is the incoming flow. The flow μ_j at a node j with two children l and r satisfies $\mu_j \geq \mu_l + \mu_r$. The amount of flow leaked at the node is $\mu_j - \mu_l - \mu_r$. In the above example, the root has flow 6 and leaks 1 unit of flow. The node having flow 1 leaks $\frac{1}{6}$. Flows generalize monotonic sequences since the flow along each path from the root to a leaf must be a nonincreasing sequence.*

By convention, μ_1 will denote the flow to the root node, and if $\mathcal{C}(i) = \emptyset$, meaning that node i is a leaf, then $\sum_{j \in \mathcal{C}(i)} \mu_j = 0$. Thus $\mu_i \geq 0$ for all $i = 1, 2 \ldots, n$.

Noisy flows can be seen as arising naturally in certain applications. For instance, suppose one seeks to estimate the population of a certain species in a large geographic region. A sampling survey may give a noisy estimate of the population μ_R in one region R, and separate estimates might be obtained in a few nonoverlapping subregions $S_1, S_2, \ldots, S_d$ with $S_i \subset R$ and $S_i \cap S_j = \emptyset$. Each subregion could be recursively partitioned. The populations are governed by the flow constraint $\mu_R \geq \mu_{S_1} + \cdots + \mu_{S_d}$. We refer to Skinner et al. (1989) for a treatment of classical survey sampling and its connection to hierarchical modeling.

As another example, computer program profiling provides a set of techniques to analyze the runtime behavior of a program, measuring the time and storage used in different parts of the code (Graham et al., 1982; Spivey, 2004). Profiling software typically generates a tree (or a more general directed graph) showing the execution time spent in different blocks of code. The flow constraints pertain since the time spent in calls to a specific function must be greater than total time spent in blocks of code that are reached from that function. Profiling is achieved by instrumenting the compiled code with instructions to monitor its performance. But instrumentation may change the performance characteristics of the program. Statistical profilers use sampling to allow the program to operate closer to its true execution behavior, with fewer side-effects. More advanced statistical estimation, such as the structured regression techniques we study here, could enable more efficient sampling schemes.

Another example arises in hierarchical classification and information retrieval with respect to a taxonomy of topics, as can occur in collaborative filtering settings. For a given search query or document to be classified, each node in the taxonomy is assigned a number, representing the relevance of the node's topic to the query. For instance, the search query might be ESPN and the topic representing sports might have children nodes basketball, baseball and football. If the relevance

of the parent (sports) to the query (ESPN) is required to be no less than the sum of the relevance values to the children nodes (basketball, baseball and football), this represents a flow constraint as defined in (1.2). The use of such constraints is called "sum-based hierarchical smoothing" by Benabbas et al. (2011). Such potential applications notwithstanding, our main motivation is that flow estimation is a natural generalization of isotonic regression, which deserves study in its own right.

We consider the problem of estimating a flow $\mu \in \mathcal{F}(T_n)$ for a given tree T_n from noisy observations

$$Y_i = \mu_i + \epsilon_i, \qquad \text{for } i = 1, \ldots, n \tag{1.3}$$

where $\mu \in \mathcal{F}(T_n)$ is unknown and the random errors $\epsilon_i \sim N(0, \sigma^2)$ are independent with σ^2 unknown. A natural estimator for μ is the least squares estimator (LSE) $\widehat{\mu}$, defined according to

$$\widehat{\mu} := \operatorname*{argmin}_{f \in \mathcal{F}(T_n)} \sum_{i=1}^{n} (Y_i - f_i)^2. \tag{1.4}$$

The LSE $\widehat{\mu}$ is uniquely defined, as it is the projection of Y onto the closed convex set $\mathcal{F}(T_n) \subset \mathbb{R}^n$. We study the behavior of $\widehat{\mu}$ under the squared error loss

$$\ell(\mu, \mu') := \frac{1}{n} \sum_{i=1}^{n} (\mu_i - \mu_i')^2. \tag{1.5}$$

The risk of any estimator $\widetilde{\mu}$ of μ under the squared loss function is given by

$$R(\mu, \widetilde{\mu}) := \mathbb{E}_\mu \ell(\widetilde{\mu}, \mu), \tag{1.6}$$

where $\mathbb{E}_\mu$ denotes the expectation taken with respect to Y having the distribution given by (1.3). In particular, $R(\mu, \widehat{\mu})$ denotes the risk of the LSE.

We study the behavior of the LSE in a setting where the number of nodes n increases in a sequence of trees T_n. The central statistical questions we investigate include the following.

- For a given sequence of trees T_n, what is the behavior of the risk of the least squares estimator $R(\widehat{\mu}, \mu)$? Is it consistent, in the sense that $R(\widehat{\mu}, \mu) \longrightarrow 0$ as the number of nodes n increases? If so, what is the rate of convergence of the risk of the LSE $R(\mu, \widehat{\mu})$? How does $R(\widehat{\mu}, \mu)$ depend on the choice of the sequence of trees?

- What is the fundamental limit of estimation in the minimax sense? In other words, what is the scaling as $n \to \infty$ of the quantity

$$\Delta_n = \inf_{\widetilde{\mu}} \sup_{\mu \in \mathcal{F}(T_n)} R(\widehat{\mu}, \mu), \tag{1.7}$$

  and how does the minimax rate of estimation depend on the structure of the underlying trees?

We provide some answers to these questions in this paper, which appears to be the first time flow estimation has been studied from a statistical perspective.

Flow estimation is a generalization of the well studied isotonic vector estimation problem. In particular, if the tree is the path graph L_n, the problem is to estimate $\mu = (\mu_1, \dots, \mu_n) \in \mathbb{R}^n$ from observations

$$Y_i = \mu_i + \epsilon_i \qquad \text{for } i = 1, \dots, n$$

under the constraint $\mu_1 \geq \cdots \geq \mu_n \geq 0$. This is, of course, a special case of univariate isotonic regression and has a long history; see e.g., Ayer et al. (1955); Brunk (1955); van Eeden (1958). The risk of the LSE for isotonic regression has been studied by a number of authors, including Birgé and Massart (1993); Chatterjee et al. (2015b); Donoho (1991); Meyer and Woodroofe (2000); van de Geer (1990, 1993); Wang (1996); Zhang (2002). It is shown by Zhang (2002) that the risk satisfies

$$R(\mu, \widehat{\mu}) \leq C \left\{ \left(\frac{\sigma^2 V(\mu)}{n} \right)^{2/3} + \frac{\sigma^2 \log n}{n} \right\}, \tag{1.8}$$

with $V(\mu) := \mu_1 - \mu_n$, where C is a universal positive constant. This result shows that the risk of $\widehat{\mu}$ scales as $n^{-2/3}$ provided $V(\mu)$ is bounded from above by a constant; this is in fact the minimax rate of estimation in this problem (see e.g., Zhang (2002)).

More broadly, flow denoising is an example of graph-based signal estimation, a topic that is of increasing recent interest. To mention a few recent results in this vein, a lasso-type penalized estimator has been proposed by Sharpnack (2013) to estimate sparse signals; Wang et al. (2014) propose adapting trend filtering ideas from nonparametric regression (Kim et al., 2009; Tibshirani et al., 2014). Such approaches are based on penalized empirical risk, and require a tuning parameter that can be difficult to set in practice. In contrast, the LSE for flow denoising requires no tuning parameters; in this way it resembles shape constrained estimation problems such as convex regression or log-concave density estimation.

The following section presents our approach to investigating the questions posed above for flow estimation on different families of trees, and states our technical results. Our main finding is a surprising gap between the rate of convergence for the least squares estimator and the minimax rate over all possible estimators, where the rate for the LSE is not in general monotonic with respect to the depth of the tree. Section 3 gives the proofs of these results. Simulations supporting our analysis are provided in Section 4.

2. Results

We analyze two main regimes, corresponding to different assumptions on the family of trees T_n. In the first regime the depth, or diameter, of the tree stays bounded or grows at most logarithmically with the number of nodes n. Our second regime of study is a family of trees containing many long paths. The depth of the trees in this regime grows at a polynomial rate.

2.1. Bounded Depth Trees

In this regime we sharply characterize the risk of the least squares estimator. Our first theorem gives an upper bound on the risk $R(\widehat{\mu}, \mu)$ in terms of the tree depth (or height) h_n, defined as the

maximum graph distance from the root to a leaf. We denote the Euclidean norm by $\|\cdot\|$, and let μ_1 denote the flow at the root for any $\mu \in \mathcal{F}(T_n)$.

Theorem 2.1. *For any tree T with height h, the worst case risk of the the least squares estimator satisfies*

$$\sup_{\mu \in \mathcal{F}(T):\mu_1 \leq V} \mathbb{E}\frac{1}{n}\|\widehat{\mu} - \mu\|^2 \leq C\frac{1}{n}\Big(\sigma^2 h^2 \log\big(\frac{n}{h}\big) + hV\sigma\sqrt{\log\big(\frac{n}{h}\big)}\Big) \tag{2.1}$$

for some universal constant C.

This risk bound holds generally, but it is mainly useful when the depth of the tree grows very slowly with n. For example, for trees with depth growing logarithmically with n, it implies an almost parametric rate of convergence $\widetilde{O}(1/n)$. Examples include the complete binary tree and the star graph with n vertices.

Restricting to the case of bounded depth trees, with height $h_n \leq C$ where C is a constant, Theorem 2.1 shows that the risk of the least squares estimator scales according to $\sigma^2 \log n/n + V\sigma\sqrt{\log n}/n$. Our next result is a minimax bound in this setting.

Theorem 2.2. *Let $\mathcal{F}_{n,V} \subset \mathbb{R}^n$ be the space of flows on trees of bounded depth $h_n \leq C$, with root flow $\mu_1 \leq V$. Then for any $\epsilon > 0$,*

$$\inf_{\widetilde{\mu}} \sup_{\mu \in \mathcal{F}_{n,V}} \mathbb{E}\frac{1}{n}\|\widetilde{\mu} - \mu\|^2 \geq \begin{cases} c\dfrac{V^2}{n} & 0 < V \leq \sigma\sqrt{\log n} \\ c_\epsilon V\sigma\dfrac{\sqrt{\log n}}{n} & \sigma\sqrt{\log n} < V \leq \sigma n^{1-\epsilon} \\ c\sigma^2 & V > n\sigma, \end{cases}$$

where c is a universal constant and c_ϵ is a constant that depends only on ϵ.

Remark 2.1. *The trivial estimator $\widehat{\mu} = Y$ incurs a risk of σ^2, and becomes minimax rate optimal as V increases. One expects that the minimax rate grows continuously with V, keeping n and σ fixed. The three cases above are a result of our proof technique, which only allows V to grow as $\sigma n^{1-\epsilon}$ for some arbitrarily small, but fixed, ϵ.*

In the bounded depth case, Theorem 2.1 and Theorem 2.2 together show that when $\sigma\sqrt{\log n} < V$, the minimax rate of estimation is indeed $V\sigma\sqrt{\log n}/n$, and that this is achieved by the LSE. However, when $V \leq \sigma\sqrt{\log n}$, the lower bound given in Theorem 2.2 is V^2/n which is attained by the trivial estimator $\widehat{\mu} = 0$. If the upper bound in Theorem 2.1 is tight, it indicates that the LSE may not be minimax rate optimal in the small V regime. The following result shows that this is indeed the case.

Theorem 2.3. *Under the same assumptions as in Theorem 2.1, the worst case risk of the LSE satisfies the lower bound*

$$\sup_{\mu \in \mathcal{F}(T_n):\mu_1 \leq V} \mathbb{E}\frac{1}{n}\|\widehat{\mu} - \mu\|^2 \geq C\frac{\sigma^2}{h}\frac{\log(n/h)}{n}.$$

Remark 2.2. *For trees with bounded depth, the lower bound is $C\sigma^2 \log(n)/n$. Theorem 2.3 actually holds for the zero flow $\mu = 0$. In other words, the risk of the LSE at the origin is $\sigma^2 \log(n)/n$, up to a constant factor. This follows from an argument involving the statistical dimension of the cone of flows for the star graph; see Amelunxen et al. (2013).*

Remark 2.3. *A simple example of a sequence of trees with bounded depth is the collection of star graphs, with one root and $n-1$ children. With known root value μ_1, the simplex constraint makes the estimation problem similar to that of estimating a vector lying in an ℓ_1 ball of known radius.*

The above theorems quantify how the flow estimation problem is "easier" in the bounded depth regime than in the setting of a path. Specifically, the risk of the LSE in the bounded depth tree setting scales according to the (nearly) parametric rate $O(\log(n)/n)$, while in the path setting it scales as $O(n^{-2/3})$. A natural question is then whether paths are the "hardest" cases for flow estimation among all rooted trees, that is, whether $O(n^{-2/3})$ is the slowest rate among all sequences of trees. Another question is whether or not the rate of convergence of the LSE is monotonic with respect to the depth. These questions motivate our study of a family of "deep trees," as described in the next section.

2.2. Deep Trees

We define a family of trees $\mathcal{T}_{\alpha,n}$ parameterized by α satisfying $0 \le \alpha \le 1$. For a given n and α, the root has $m = \lceil n^\alpha \rceil$ children, and each of these m children is the starting point of a path of length $l = \lceil n^{1-\alpha} \rceil$. See Figure 2. When $\alpha = 0$ the tree $\mathcal{T}_{\alpha,n}$ is a single path and flow denoising on $\mathcal{T}_{\alpha,n}$ corresponds to the isotonic sequence regression problem. When $\alpha = 1$ the tree $\mathcal{T}_{\alpha,n}$ is a star graph with n children of the root. Hence, this family of trees interpolates between the path and star graphs. If the risk of the LSE is indeed monotonic with the depth of the underlying sequence of trees, then it should decay faster with n as α increases.

2.2.1. Upper Bounds for the LSE

When α is closer to zero than to one, a natural estimator is to set $\widehat{\mu}_1 = Y_1$ at the root and perform n^α separate isotonic regressions on each of the paths. Since this estimator need not satisfy the flow constraint at the root, we then project the resulting estimator onto the space $\mathcal{F}(\mathcal{T}_{n,\alpha})$ to obtain the estimator $\widehat{\mu}_{\text{nat}}$.

Using the results of Zhang (2002) for isotonic regression, together with an application of Hölder's inequality (see Section 3), we can show that the risk of the natural estimator satisfies

$$\sup_{\mu \in \mathcal{F}(\mathcal{T}_{\alpha,n}) : \mu_1 \le V} \mathbb{E} \frac{1}{n} \|\widehat{\mu}_{\text{nat}} - \mu\|^2 \le \left(\frac{\sigma^2 V}{n}\right)^{2/3} + \frac{\sigma^2 \log en}{n^{1-\alpha}}. \tag{2.2}$$

For small α, it is natural to expect an "isotonic effect" where the performance of the actual LSE $\widehat{\mu}$ should be close to the performance of the natural estimator estimator $\widehat{\mu}_{\text{nat}}$. This intuition is made rigorous by our first upper bound for the LSE over $\mathcal{F}(\mathcal{T}_{\alpha,n})$.

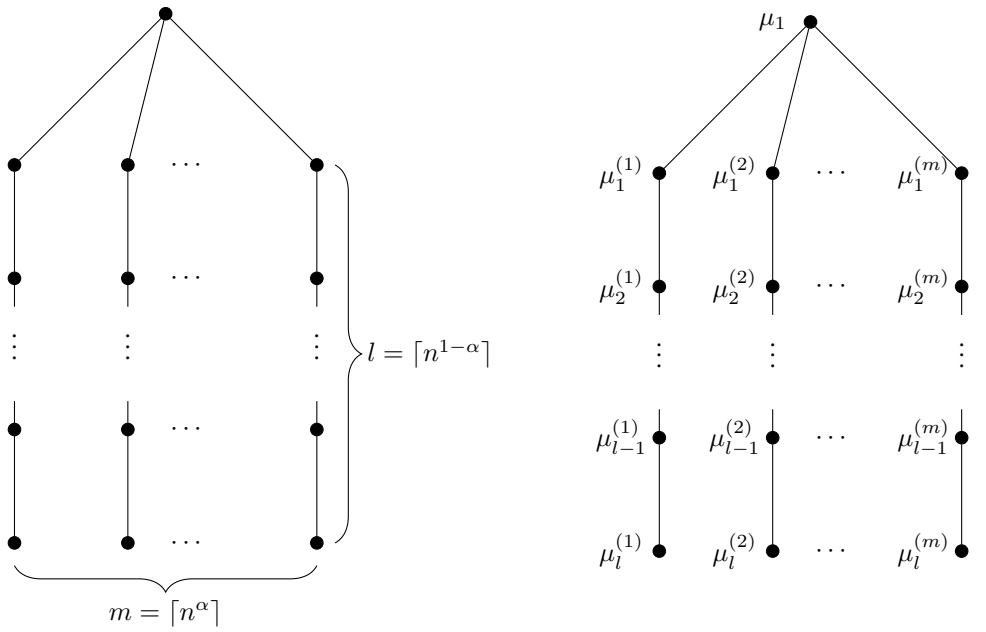

FIG 2. *Left: A family of trees $\mathcal{T}_{\alpha,n}$, consisting of root with $m = \lceil n^{\alpha} \rceil$ children, and m paths of length $l = \lceil n^{1-\alpha} \rceil$. Right: A parameterization of flows in the family $\mathcal{F}_{\alpha,n} = \mathcal{F}(\mathcal{T}_{\alpha,n})$.*

Theorem 2.4. *The risk of the LSE satisfies*

$$\sup_{\mu \in \mathcal{F}(\mathcal{T}_{\alpha,n}) : \mu_1 \le V} \mathbb{E} \frac{1}{n} \|\widehat{\mu} - \mu\|^2 \le 2 \Big(\frac{\sigma^2 V \log(en)}{n} \Big)^{2/3} + \frac{\sigma^2 \log(en)}{n^{1-\alpha}}.$$

The proof of Theorem 2.4 exploits monotonicity of the flow along the paths, but does not use the simplex constraints at every level. The proof technique is based on a general theory of analyzing least squares estimators which uses an appropriate notion of size of the tangent cone at flows that are piecewise constant along every path. This point will be made clear in Section 3, where Theorem 2.4 is proved.

As α increases, one expects that imposing the simplex constraint at the root becomes more important. Since Theorem 2.4 is proved only using monotonicity along each path, a separate risk bound is needed in the large α regime.

This becomes clear if one thinks about the zero estimator $\widehat{\mu} = 0$. The sum of the squared errors of this estimator is V^2 at every level of the tree; since the tree has depth $n^{1-\alpha}$, the sum of squared errors of the zero estimator is at most $n^{1-\alpha} V^2$. The resulting bound can be better than the upper bound given in Theorem 2.4.

As a type of oracle estimator that imposes the flow constraints, suppose we know the value of $\mu_1 = V$. Then we could perform "simplex regression" on each level. Specifically, using the notation of Figure 2, we first compute, for each level ℓ, the estimator

$$\widetilde{\mu}_\ell = \operatorname*{argmin}_{\sum_{j=1}^m \mu_\ell^{(j)} \le V} \sum_{j=1}^m \big(Y_\ell^{(j)} - \mu_\ell^{(j)} \big)^2 . \tag{2.3}$$

We then project $\widetilde{\mu}$ onto the set of flows. Let us call this estimator $\widehat{\mu}_{oracle}$, since it assumes knowledge of the root flow V. We can then establish the risk bound

$$\mathbb{E}\frac{1}{n}\|\widehat{\mu}_{oracle} - \mu\|^2 \leq C\left(V\sigma\frac{\sqrt{\log n}}{n^\alpha} + \sigma^2\frac{\log n}{n^\alpha}\right) \tag{2.4}$$

using Theorem 2.1. For large α, one expects that the performance of the LSE will be similar to this oracle estimator. Our next result establishes that this is indeed the case.

Theorem 2.5. *The risk of the LSE satisfies*

$$\mathbb{E}\frac{1}{n}\|\widehat{\mu} - \mu\|^2 \leq C\big(\frac{\sigma^2(\log(en))^3}{n^\alpha} + V\frac{\sigma(\log(en))^{3/2}}{n^\alpha}\big)$$

for a universal constant C.

Theorem 2.5 is proved using techniques that control the maxima of a suitable empirical process, requiring the estimation of covering numbers for the space of flows. The full proof is given in Section 3.5. We remark that the extra log factor in the bound in Theorem 2.5 appears to be an artifact of our current proof technique, and may not be necessary.

Theorems 2.4 and 2.5 hold for any $0 \leq \alpha \leq 1$, but are useful in different regimes. The following corollary combines them to establish a single risk bound. This is done by treating V and σ as fixed, and considering the risk as a function of n only. The scaling of n is better in Theorem 2.4 for $0 \leq \alpha \leq 1/3$ and is better in Theorem 2.5 when $1/3 < \alpha \leq 1$.

Corollary 2.6. *The worst case risk of the least squares estimator satisfies*

$$\sup_{\mu\in\mathcal{F}(\mathcal{T}_{n,\alpha}):\mu_1\leq V}\mathbb{E}\frac{1}{n}\|\widehat{\mu} - \mu\|^2 \leq \begin{cases} 2\big(\dfrac{\sigma^2 V\log(en)}{n}\big)^{2/3} + \dfrac{\sigma^2\log(en)}{n^{1-\alpha}} & 0 < \alpha \leq \frac{1}{3} \\ \dfrac{C}{n}\big(\sigma^2(\log(en))^3 n^{1-\alpha} + V\sigma(\log(en))^{3/2}\big) & \frac{1}{3} \leq \alpha < 1. \end{cases}$$

Note that for fixed V and σ, the upper bound of Corollary 2.6 is not monotonically decreasing in α. To resolve whether this is intrinsic to the flow estimation problem, lower bounds are required, as discussed in the next section.

2.2.2. Minimax Lower Bounds

Lower bounds for flow estimation in the family of trees $\mathcal{T}_{n,\alpha}$ should match the lower bound for isotonic regression when $\alpha = 0$. We therefore first establish a minimax lower bound for isotonic regression; although existing lower bounds are comparable (Chatterjee et al., 2015b), we are not aware of this exact form of the lower bound appearing in the literature.

Theorem 2.7. *The minimax lower bound*

$$\inf_{\widetilde{\mu}}\sup_{\mu\in\mathcal{M}_V}\mathbb{E}\frac{1}{n}\|\widetilde{\mu} - \mu\|^2 \geq \frac{1}{192}\min\Big\{\sigma^2,\, V^2,\, \Big(\frac{\sigma^2 V}{n}\Big)^{2/3}\Big\}$$

holds for the parameter space

$$\mathcal{M}_V = \{\mu \in \mathbb{R}^n : V \geq \mu_1 \geq \cdots \geq \mu_n \geq 0\}.$$

This bound is a minimum of three terms. If V is not too small or large, the rate $\big(\sigma^2 V/n\big)^{2/3}$ is standard and is attained by the LSE. If V is very small, the V^2 term is the smallest, and this rate is achieved by the trivial estimator $\widehat{\mu} = 0$. In this regime, the LSE is not minimax rate optimal since it suffers a risk of at least $\sigma^2 \log(en)/n$ at the origin; see Remark 2.2. In case V is very large, the rate becomes σ^2, which is achieved by the trivial estimator $\widehat{\mu} = Y$.

For α small, the hardness of flow estimation on $\mathcal{T}_{n,\alpha}$ is dominated by the monotonicity constraints in each path rather than by the simplex constraints. This leads us to derive a minimax lower bound using Theorem 2.7.

Theorem 2.8. *Let $\mathcal{F}_{n,\alpha,V}$ denote the set of flows on tree $\mathcal{T}_{n,\alpha}$ with root value at most V, and set $m = \lceil n^\alpha \rceil$. Then for any α,*

$$\inf_{\widetilde{\mu}} \sup_{\mu \in \mathcal{F}_{n,\alpha,V}} \mathbb{E} \frac{1}{n} \|\widetilde{\mu} - \mu\|^2 \geq \sup_{(v_1,\ldots,v_m):\, \sum_{i=1}^m v_i \leq V} \frac{c}{n} \sum_{i=1}^m \min\Big\{ n^{1-\alpha}\sigma^2,\ (v_i\sigma^2)^{2/3} n^{(1-\alpha)/3},\ n^{1-\alpha} v_i^2 \Big\}.$$

This result is proved by considering the subset of flows where the m the children of the root have flows $(v_1, \ldots, v_m)$, together satisfying $\sum_{i=1}^m v_i \leq V$. Estimation in this parameter space is then equivalent to m separate isotonic regressions and the minimax lower bound is obtained by using Theorem 2.7. We obtain, as a corollary, a minimax lower bound in the regime where α is small and V is neither too small nor too large for the trivial estimators to dominate.

Corollary 2.9. *Let $0 \leq \alpha \leq 1/3$, and suppose that $n^{3\alpha - 1} \leq V \leq n$. Then*

$$\inf_{\widetilde{\mu}} \sup_{\mu \in \mathcal{F}_{n,\alpha,V}} \mathbb{E} \frac{1}{n} \|\widetilde{\mu} - \mu\|^2 \geq c \Big(\frac{V\sigma^2}{n} \Big)^{2/3}.$$

Thus, under the assumptions of this corollary, the LSE is minimax rate optimal up to a constant factor.

For the range $1/3 \leq \alpha \leq 1$, we prove lower bounds in a different manner, using a similar strategy as was used in the lower bound for the star graph.

Theorem 2.10. *For fixed $1/3 < \alpha \leq 1$,*

$$\inf_{\widetilde{\mu}} \sup_{\mu \in \mathcal{F}_{n,\alpha,V}} \mathbb{E} \frac{1}{n} \|\widetilde{\mu} - \mu\|^2 \geq \begin{cases} c \dfrac{V^2}{n^\alpha} & \text{in case } 0 < V \leq \dfrac{\sigma\sqrt{\log n}}{n^{(1-\alpha)/2}} \\[2ex] c_\epsilon V \sigma \dfrac{\sqrt{\log n}}{n^{(1+\alpha)/2}} & \text{for } \dfrac{\sigma\sqrt{\log n}}{n^{(1-\alpha)/2}} < V \leq n^{(1+\alpha)/2 - \epsilon}\sigma, \end{cases}$$

for any $\epsilon > 0$, where c is a universal constant and c_ϵ is a constant only depending on ϵ.

In the first case of this result, the trivial estimator $\widehat{\mu} = 0$ attains the optimal rate. However, in the second case, where V is large, the lower bound does not match our bound for the LSE in Theorem 2.5. Our next result shows that, in fact, our minimax lower bound is tight (up to log factors) in this regime.

Theorem 2.11. *For any $1/3 \leq \alpha \leq 1$, we have the upper bound*

$$\inf_{\widetilde{\mu}} \sup_{\mu \in \mathcal{F}_{n,\alpha,V}} \mathbb{E} \frac{1}{n} \|\widetilde{\mu} - \mu\|^2 \leq 21 V \sigma \frac{\sqrt{\log(en)}}{n^{(1+\alpha)/2}} + 16\sigma^2 \frac{1 + \log(en)}{n}. \tag{2.5}$$

This bound is established by considering a least squares estimator on an appropriate finite net over the space of flows $\mathcal{F}(\mathcal{T}_{n,\alpha})$. The proof is thus an information theoretic argument, and does not exhibit a computationally efficient estimator.

Remark 2.4. *As a possible avenue for constructing a computationally efficient estimator, suppose that a procedure can detect which of the children of the root have small flow value. If the entire path under such nodes is estimated as zero, and otherwise an isotonic regression is fit, this might attain the minimax rate. We leave this as a topic for future research.*

2.3. Lower Bound for the LSE

Finally, we address the question of whether the LSE is rate optimal for the range $1/3 < \alpha \leq 1$, noting that the $\widetilde{O}(n^{1-\alpha})$ upper bound of Theorem 2.5 does not match the $\widetilde{\Omega}(n^{(1-\alpha)/2})$ lower bound of Theorem 2.10. The following result asserts that there is an actual gap.

Theorem 2.12. *For any $V > 0$, the worst case risk of the least squares estimator satisfies the lower bound*

$$\sup_{\mu \in \mathcal{F}_{n,\alpha,V}} \mathbb{E} \frac{1}{n} \|\widehat{\mu} - \mu\|^2 \geq \begin{cases} \pi^2 \sigma^2 n^{\alpha - 1} & 1/3 \leq \alpha \leq 1/2 \\ \pi^2 \sigma^2 n^{-\alpha} & 1/2 \leq \alpha \leq 1. \end{cases}$$

The lower bound in Theorem 2.12 actually holds when $\mu = 0$ is the zero flow; the proof of this result is given in Section 3.10.

Taken together, all of these bounds characterize the exponents in the worst case risk for the LSE and any possible flow estimator, as portrayed graphically in Figure 3. They imply the perhaps surprising conclusion that the LSE is not rate optimal over the range $1/3 \leq \alpha \leq 1$, considering V and σ to be fixed constants.

3. Proofs

In this section we give the proofs of the results described in the previous section. These results characterize the performance of the LSE for flows and the minimax rates of convergence, and reveal a gap between the LSE and the best possible estimators for flows on trees. In order to get sharp results, we employ a variety of proof techniques, as shown in Table 1.

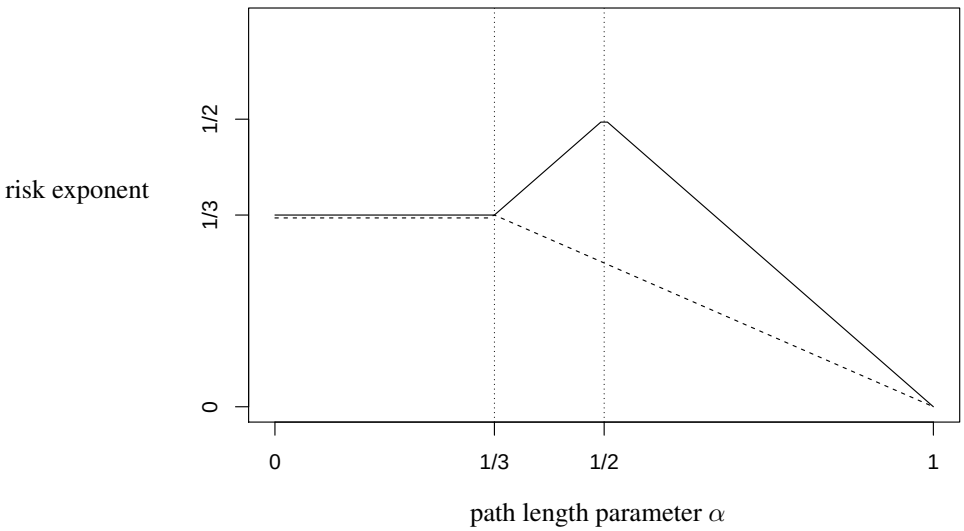

FIG 3. *With $\alpha = 0$ the tree $T_{n,\alpha}$ is a single path, and with $\alpha = 1$ the graph is a star; in general the tree has $m = \lceil n^{\alpha} \rceil$ paths, each of length $l = \lceil n^{1-\alpha} \rceil$. This plot shows the leading exponent (thus disregarding log factors) in the expected squared error $\mathbb{E}\|\widehat{\mu} - \mu\|^2$ (which is n times the risk) of the LSE (solid line) the minimax lower bound (dashed line). Our sequence of results establishes that these exponents are tight, and thus exhibit an unusual gap in a shape-constrained estimation problem between the performance of the least squares estimator and the best possible estimator.*

3.1. Proof of Theorem 2.1: Upper bound on LSE for shallow trees

A key result we use to prove our risk bounds for the LSE is the following technique based on a quadratic Gaussian supremum functional, due to Sourav Chatterjee (Chatterjee, 2014b).

Theorem 3.1 (Chatterjee, 2014)**.** *Fix a rooted tree T and a flow $\mu \in \mathcal{F}(T)$. Define the function $f_\mu : \mathbb{R}_+ \to \mathbb{R}$ as*

$$f_\mu(t) := \mathbb{E}\left(\sup_{\nu \in \mathcal{F} : \|\nu - \mu\| \le t} \langle Z, \nu - \mu \rangle \right) - \frac{t^2}{2} \tag{3.1}$$

where Z is an n-dimensional mean zero Gaussian vector with covariance matrix $\sigma^2 I$. Let $t^* > 0$ satisfy $f_\mu(t^*) \le 0$. Then there exists a universal positive constant C such that

$$\mathbb{E}\frac{1}{n}\|\widehat{\mu} - \mu\|^2 \le \frac{C}{n} \max\left(t^{*2}, \sigma^2\right). \tag{3.2}$$

The above theorem is a consequence of Chatterjee (2014b, Theorem 1.1) and Chatterjee (2014b, Proposition 1.3). Let t_μ be the point in $[0, \infty)$ where $t \mapsto f_\mu(t)$ attains its maximum; existence and uniqueness of t_μ are proved in Chatterjee (2014b, Theorem 1.1). Then there exists a universal constant C such that

$$\mathbb{E}\frac{1}{n}\|\widehat{\mu} - \mu\|^2 \le \frac{C}{n} \max\left(t_\mu^2, \sigma^2\right).$$

This reduces the problem of bounding $\mathbb{E}\frac{1}{n}\|\widehat{\mu} - \mu\|^2$ to that of bounding t_μ. For this latter problem, Chatterjee (2014b, Proposition 1.3) observes that

$$t_\mu \le t^* \text{ whenever } t^* > 0 \text{ and } f_\mu(t^*) \le 0.$$

	Result	*Proof technique*
Theorem 2.1	Upper bound for LSE; shallow trees	Gaussian supremum functional
Theorem 2.2	Minimax lower bound; shallow trees	Fano's lemma
Theorem 2.3	Lower bound for LSE; shallow trees	Gaussian widths
Theorem 2.4	Isotonic upper bound for LSE; deep trees	Statistical dimension
Theorem 2.5	Simplex upper bound for LSE; deep trees	Chaining and entropy bounds
Theorem 2.7	Minimax lower bound; monotone sequences	Assouad's lemma
Theorem 2.8	Minimax lower bound; deep trees, $\alpha \leq \frac{1}{3}$	Minimax for isotonic regression
Theorem 2.10	Minimax lower bound; deep trees, $\alpha \geq \frac{1}{3}$	Fano's lemma
Theorem 2.11	Tightness of minimax lower bound, $\alpha \geq \frac{1}{3}$	Covering numbers, LSE on a net
Theorem 2.12	Tightness of LSE upper bound, $\alpha \geq \frac{1}{3}$	Gaussian widths

TABLE 1
A catalogue of the results, with the main techniques used in their proofs.

In order to bound t_μ, one therefore seeks $t^* > 0$ such that $f_\mu(t^*) \leq 0$. This now requires a bound on the expected supremum of the Gaussian process in the definition of $f_\mu(t)$ in (3.1).

Proof of Theorem 2.1. Fix any $\mu \in \mathcal{F}(T_n)$. Let us denote by $g_\mu(t)$ the supremum

$$g_\mu(t) = \sup_{\nu \in \mathcal{F}(T_n): \|\nu - \mu\| \leq t} \langle Z, \nu - \mu \rangle.$$

Denote by L_i the set of nodes in level $1 \leq i \leq h$ in the tree T_n. Then we can write

$$\begin{aligned} g_\mu(t) &\leq \sup_{\nu \in \mathcal{F}(T_n): \|\nu - \mu\| \leq t} Z_1(\nu_1 - \mu_1) + \sum_{i=2}^{h} \sup_{\nu \in \mathcal{F}(T_n): \|\nu - \mu\| \leq t} \sum_{j \in L_i} Z_j(\nu_j - \mu_j) \\ &\leq t|Z_1| + \sum_{i=2}^{h} \sup_{\nu \in \mathcal{F}(T_n): \|\nu - \mu\| \leq t} \Big(\sum_{j \in L_i} Z_j \nu_j + \sum_{j \in L_i} (-Z_j) \mu_j \Big). \end{aligned}$$

The last inequality is because the constraint $\|\nu - \mu\| \leq t$ implies that $\nu_1 \leq \mu_1 + t$. The flow constraints imply, within each level L_i, that the sum $\sum_{j \in L_i} \nu_j$ is at most $\mu_1 + t$. Therefore, the last equation can be further bounded to obtain

$$g_\mu(t) \leq t|Z_1| + \sum_{i=2}^{h} \big(\max_{j \in L_i} Z_j\big)(\mu_1 + t) + \big(\max_{j \in L_i} -Z_j\big)\mu_1$$

where we have used the fact that $\sum_{i=1}^d a_i b_i \leq (\sum_{i=1}^d b_i) \max_{1 \leq i \leq d} a_i$ for any arbitrary vector a and any nonnegative vector b.

Now taking the expectation of both sides, and using the fact that the expectation of the maximum of d independent $N(0, \sigma^2)$ random variables is upper bounded by $\sigma\sqrt{2 \log d}$ and $E|Z_1| = \sigma\sqrt{\frac{2}{\pi}}$,

we obtain

$$\begin{aligned}\mathbb{E}g_\mu(t) &\le t\sigma\sqrt{\frac{2}{\pi}} + \sum_{i=2}^{h}(2\mu_1 + t)\sigma\sqrt{2\log|L_i|} \\ &\le t\sigma\sqrt{\frac{2}{\pi}} + h(2\mu_1 + t)\sigma\sqrt{2\log\frac{n}{h}},\end{aligned}$$

where we have used Jensen's inequality in the second. Recalling that $f_\mu(t) = \mathbb{E}g_\mu(t) - t^2/2$, this implies that $f_\mu(t^*) \le 0$ where

$$t^* = C\max\left\{\sigma h\sqrt{\log\frac{n}{h}},\ \sqrt{\mu_1\sigma h\sqrt{\log\frac{n}{h}}}\right\}.$$

Theorem 3.1 now finishes the proof of the upper bound. □

3.2. Proof of Theorem 2.2: Minimax lower bound for shallow trees

To prove this minimax lower bound we use the following standard version of Fano's lemma. For any set $\mathcal{A} \subset \mathbb{R}^n$ define a δ-packing set of $\mathcal{A}$ to be a finite subset of $\mathcal{A}$ such that for any two distinct points $a \ne a' \in \mathcal{A}$ we have separation $\|a - a'\|^2 \ge \delta$.

Lemma 3.2 (Fano)**.** *If W is a δ-packing set of Θ then*

$$\inf_{\widehat{\theta}}\sup_{\theta\in\Theta}\mathbb{E}\|\widehat{\theta} - \theta\|^2 \ge \frac{\delta}{2}\Big(1 - \frac{\Delta + \log 2}{\log|W|}\Big)$$

where $\Delta = \max_{w\ne w'} KL(\mathbb{P}_w, \mathbb{P}_{w'})$, with $\mathbb{P}_w$ denoting the distribution of the data Y when the true underlying parameter is w.

We also require the following version of the Varshamov-Gilbert coding lemma.

Lemma 3.3 (Varshamov-Gilbert)**.** *Let k and d be two positive integers with $k \le d/6$. Then there exists a set $W \in \{0,1\}^d$ such that the following three conditions hold:*

1. *For any $w \in W$ we have $\sum_{i=1}^d w_i = k$.*
2. *For any $w \ne w' \in W$ we have $\sum_{i=1}^d \mathcal{I}\{w_i \ne w'_i\} \ge k/2$.*
3. *$\log|W| \ge \frac{k}{2}\log(1 + \frac{d}{4k})$.*

Proof of Theorem 2.2. We first prove the minimax lower bound for the star graph with one root and n children, with the parameter space $\mathcal{F}_{n,V} \subset \mathbb{R}^n$ being the space of flows on the star graph with root value at most V. For any fixed $k \le n/2$ consider the packing set W guaranteed by Lemma 3.3. For each element $w \in W$, we obtain a flow on the star graph by setting the root flow to be V, and setting the children flows to $V/k \cdot w \in \mathbb{R}^n_+$. This forms a packing set of $\mathcal{F}_{n,V}$ of squared radius

$V^2/2k$. The Euclidean distance between any two distinct elements in this packing set is at most $2V^2/k$. Therefore an application of Lemma 3.2 gives the minimax lower bound

$$\inf_{\widetilde{\theta}} \sup_{\theta \in \mathcal{F}_{n,V}} \mathbb{E}\|\widetilde{\theta} - \theta\|^2 \geq \max_{1 \leq k \leq n/6} \left\{ \frac{V^2}{8k} \Big(1 - \frac{V^2/(2k\sigma^2) + \log 2}{\frac{k}{2}\big(\log(1 + \frac{n}{4k}) \big)} \Big) \right\}.$$

By choosing k so that $1 \leq k \leq n/60$ we can ensure

$$\frac{\log 2}{\frac{k}{2}\big(\log(1 + \frac{n}{4k}) \big)} \leq \frac{2 \log 2}{\log(16)} \leq \frac{1}{2}.$$

and thus obtain the somewhat cleaner bound

$$\inf_{\widetilde{\theta}} \sup_{\theta \in \S_V} \mathbb{E}\|\widetilde{\theta} - \theta\|^2 \geq \max_{1 \leq k \leq n/60} \left\{ \frac{V^2}{8k} \Big(\frac{1}{2} - \frac{V^2}{k^2\sigma^2\big(\log(1 + \frac{n}{4k}) \big)} \Big) \right\}. \tag{3.3}$$

We now select k differently according to the scale of the root flow V.

Small flow: $0 \leq V \leq \sigma\sqrt{\log n}$**.** In this case we set $k = c$ for some appropriate constant $c \geq 1$. Then we have

$$\frac{V^2}{k^2\sigma^2\big(\log(1 + \frac{n}{4k}) \big)} = \frac{V^2}{c^2\sigma^2\big(\log(1 + \frac{n}{4c}) \big)} \leq \frac{\sigma^2 \log n}{c^2\sigma^2\big(\log(1 + \frac{n}{4c}) \big)} \leq \frac{1}{4}$$

where the first inequality uses the bound on V and the last inequality holds for all large enough n if c is chosen to be a sufficiently large constant. Together with (3.3), this implies a minimax lower bound of V^2/c.

Large flow: $\sigma\sqrt{\log n} < V \leq n^{1-\epsilon}\sigma$**.** Let $c_\epsilon \geq 1$ be a sufficiently large positive constant. Set $k = c_\epsilon \frac{V}{\sigma\sqrt{\log n}}$. Then $k \geq 1$ by the lower bound condition on V. By the upper bound condition, for sufficiently large n we have $k \leq c_\epsilon \frac{n\sigma}{\sigma \log n} \leq \frac{n}{60}$. Hence the choice of k is valid for all large enough n, and we have

$$\frac{V^2}{k^2\sigma^2\big(\log(1 + \frac{n}{4k}) \big)} = \frac{V^2\sigma^2 \log n}{c_\epsilon V^2\sigma^2\big(\log(1 + \frac{n\sigma\sqrt{\log n}}{4c_\epsilon V}) \big)} \leq \frac{\log n}{c_\epsilon\big(\log(1 + n^\epsilon \log n) \big)} \leq \frac{1}{4}$$

where the first inequality uses the upper bound on V and the last inequality is true for all large enough n. Using (3.3) then yields the lower bound of $\frac{V^2}{k} = V\sigma\sqrt{\log n}$, up to a constant factor depending on ϵ. Note that our choice of k need not be a integer; choosing the nearest integer to it which will only change the lower bound by constants.

The proof for a general tree T with height h bounded by a constant C is very similar. In particular, there must exist a level l of the tree with at least n/C elements. Define a flow where the values at level l assume the values given by a vector $w \in W \subset \{0, 1\}^{n/C}$ multiplied by V/k. Define the flow at levels $l' < l$ by defining the value at any vertex to be the sum of the values at its children,

and the flow at levels $l' > l$ to be zero. This defines a flow indexed by an element $w \in W$. Since the height of T is at most a constant C, one can verify that this defines a packing set of $\mathcal{F}_{n,V}$ of squared radius V^2/k up to a constant factor. Also, for any two distinct elements in this packing set, the Euclidean distance between them is at most V^2/k up to a constant factor. The remainder of the argument using Fano's lemma proceeds as above. □

3.3. Proof of Theorem 2.3: Lower bound for the LSE on shallow trees

We will need the following standard lemma about the expectation of the maxima of independent normal random variables.

Lemma 3.4. *Let $Z_1, \dots, Z_n$ be independent $N(0, \sigma^2)$ random variables. Then*

$$\frac{\sqrt{2\sigma^2 \log n}}{\sqrt{\pi \log 2}} \leq \mathbb{E} \max_{1 \leq i \leq n} Z_i \leq \sqrt{2\sigma^2 \log n}.$$

We also need the following lemma about the projection to a closed convex cone.

Lemma 3.5. *Let $K \subset \mathbb{R}^n$ be a closed convex cone, and denote by $\Pi_K(z)$ the projection of $z \in \mathbb{R}^n$ to K. Then*

$$\|\Pi_K(z)\| = \sup_{\nu \in K : \|\nu\| \leq 1} \langle z, \nu \rangle. \tag{3.4}$$

Proof. By definition, $\Pi_K(z)$ maximizes $\langle z, v \rangle - \frac{1}{2}\|v\|^2$ among all $v \in K$; thus

$$\begin{aligned} \|\Pi_K(z)\| &= \operatorname*{argmax}_{t \geq 0} \left\{ \sup_{\theta \in K : \|\theta\| \leq t} \langle z, \theta \rangle - \frac{t^2}{2} \right\} \\ &= \operatorname*{argmax}_{t \geq 0} \left\{ t \sup_{\theta \in K : \|\theta\| \leq 1} \langle z, \theta \rangle - \frac{t^2}{2} \right\}, \end{aligned}$$

where the second equation uses the fact that K is a cone. Computing the maximum of the quadratic finishes the proof of the lemma. □

Proof of Theorem 2.3. We will prove the desired lower bound on the risk of the LSE at the origin, where $\mu_i = 0$ for all $1 \leq i \leq n$. Since the space of flows $\mathcal{F}(T_n)$ is a cone, Lemma 3.5 implies that

$$\|\widehat{\mu} - \mu\| = \sup_{\nu \in \mathcal{F}(T_n) : \|\nu\| \leq 1} \langle Z, \nu \rangle.$$

Therefore

$$\left(\mathbb{E} \sup_{\nu \in \mathcal{F}(T_n) : \|\nu\| \leq 1} \langle Z, \nu \rangle \right)^2 \leq \mathbb{E} \left(\sup_{\nu \in \mathcal{F}(T_n) : \|\nu\| \leq 1} \langle Z, \nu \rangle \right)^2 = \mathbb{E}\|\widehat{\mu} - \mu\|^2. \tag{3.5}$$

Thus, it suffices to lower bound the Gaussian width term $\mathbb{E} \sup_{\nu \in \mathcal{F}(T_n) : \|\nu\| \leq 1} \langle Z, \nu \rangle$. Since the tree T_n has height h, there exists a level l with no fewer than n/h vertices. Let us denote this set of vertices by L, and define

$$i^* = \operatorname*{argmax}_{i \in L} Z_i.$$

There is a unique path from the root to i^*; define $\widehat{\nu}$ to be equal to $1/\sqrt{h_n}$ on this path and equal to zero off the path. Then $\widehat{\nu} \in \mathcal{F}(T_n)$ with $\|\widehat{\nu}\| \leq 1$. Therefore we can write

$$\sup_{\nu \in \mathcal{F}(T_n):\|\nu\|\leq 1} \langle Z, \nu \rangle \geq \langle Z, \widehat{\nu} \rangle = \frac{1}{\sqrt{h_n}} \max_{i \in L} Z_i + \frac{1}{\sqrt{h_n}} \sum_{i=1}^{l-1} Z_i$$

where Z_i are the error vector coordinates as we traverse from the root to the vertex i^*. Taking expectation and applying Lemma 3.4, we have

$$\mathbb{E} \sup_{\nu \in \mathcal{F}(T_n):\|\nu\|\leq 1} \langle Z, \nu \rangle \geq \mathbb{E}\langle Z, \widehat{\nu} \rangle = \frac{1}{\sqrt{h_n}} \mathbb{E} \max_{i \in L} Z_i \geq C \frac{\sigma}{\sqrt{h_n}} \sqrt{\log \frac{n}{h_n}}.$$

This inequality combined with (3.5) finishes the proof of the lower bound. □

3.4. Proof of Theorem 2.4: Isotonic upper bound on the LSE for deep trees

We first set up some notation. Recall that for a given parameter $0 \leq \alpha \leq 1$ the root of the tree $\mathcal{T}_{\alpha,n}$ has $m = \lceil n^\alpha \rceil$ children. Each of these m children is the starting point of a path of length $l = \lceil n^{1-\alpha} \rceil$. Clearly there are $lm + 1 = \Theta(n)$ vertices in $\mathcal{T}_{\alpha,n}$. For simplicity, we will assume $lm + 1 = n$. The set of flows on $\mathcal{T}_{\alpha,n}$ is denoted by $\mathcal{F}(\mathcal{T}_{\alpha,n}) := \mathcal{F}_{n,\alpha}$ for notational simplicity. The set $\mathcal{F}_{n,\alpha}$ is a closed convex cone of $\mathbb{R}^n$. For convenience, we will index the components of a flow μ in $\mathcal{F}_{n,\alpha}$ as shown in Figure 2, with the root flow denoted by μ_1 and the values $\mu_1^{(1)}, \mu_1^{(2)}, \ldots, \mu_1^{(m)}$ denoting the flows to the m children of the root; thus $\mu_1 \geq \mu_1^{(1)} + \mu_1^{(2)} + \cdots \mu_1^{(m)}$. The monotonic nondecreasing flow along the jth path is then $\mu_1^{(j)} \geq \mu_2^{(j)} \geq \cdots \mu_l^{(j)}$. Sometimes we will denote the vector $(\mu_1^{(j)}, \mu_2^{(j)}, \mu_l^{(j)})$ by $\mu^{(j)}$. For any vector $\theta \in \mathbb{R}^n$ let $k(\theta)$ denote the cardinality of the set $\{\theta_1, \ldots, \theta_n\}$.

We will need the following lemma about approximating a monotone sequence by a piecewise constant sequence.

Lemma 3.6 (Approximation)**.** *Let $\epsilon > 0$. Fix any nonincreasing sequence $\theta \in \mathbb{R}^n$. Let $V = \theta_1 - \theta_n$. Then there exists a nonincreasing sequence $\alpha \in \mathbb{R}^n$ with $\alpha_1 \leq \theta_1$ such that $k(\alpha) \leq \frac{V}{\epsilon} + 1$ and $\|\theta - \alpha\|^2 \leq n\epsilon^2$.*

Proof. Define $s_0 = 1$. For $1 \leq i \leq \lfloor \frac{V}{\epsilon} \rfloor$ recursively define

$$s_i = \min\big\{i : \theta_{s_{i-1}} - \theta_i > \epsilon\big\}.$$

Define $s_{\lfloor \frac{V}{\epsilon} \rfloor + 1} = n$. Now define α as follows:

$$\alpha_j = \sum_{i=1}^{\lfloor \frac{V}{\epsilon} \rfloor + 1} \theta_{s_{i-1}} \mathcal{I}\{s_{i-1} \leq j < s_i\}.$$

It is clear that $\alpha_n \leq \theta_n$ and $k(\alpha) \leq \frac{V}{\epsilon} + 1$. Also by definition, we have $|\alpha_j - \theta_j| \leq \epsilon$ for all $1 \leq j \leq n$ which finishes the proof of the lemma. □

For any closed convex cone $C \subset \mathbb{R}^n$, denote the projection operator onto C by Π_C and define

$$\delta(C) = \mathbb{E}\|\Pi_C Z\|^2 \tag{3.6}$$

where $Z \sim N(0, I)$. The quantity $\delta(C)$ is called the statistical dimension of C and generalizes the concept of dimension of a subspace (see Amelunxen et al. (2013)). The tangent cone to C at $u \in \mathbb{R}^n$ is defined by

$$\mathcal{T}_{C,u} = \{v - tu : v \in C,\, t \geq 0\}. \tag{3.7}$$

The tangent cone $\mathcal{T}_{C,u}$ is a closed convex cone of $\mathbb{R}^n$ and hence one can talk about the statistical dimension $\delta(\mathcal{T}_{C,u})$.

The following oracle risk bound is due to Bellec (2015).

Lemma 3.7 (Bellec, Proposition 2.1)**.** *Let $C \subset \mathbb{R}^n$ be a closed convex cone, and let $\widehat{\mu}$ denote the least squares estimate of $\mu \in C$, that is, $\widehat{\mu} = \Pi_C(\mu + Z)$ where $Z \sim N(0, \sigma^2 I)$. Then we have the pointwise inequality*

$$\|\widehat{\mu} - \mu\|^2 \leq \|u - \mu\|^2 + \sigma^2 \left\|\Pi_{\mathcal{T}_{C,u}}(Z/\sigma)\right\|^2$$

for any $u \in C$. As a consequence,

$$\mathbb{E}\|\widehat{\mu} - \mu\|^2 \leq \|u - \mu\|^2 + \sigma^2 \delta(\mathcal{T}_{C,u}).$$

To use this result we need to bound the statistical dimension of the tangent cone to the space of flows.

Lemma 3.8. *Fix any $\nu \in \mathcal{F}_{n,\alpha}$, and let $k^{(i)} = k(\nu^{(i)})$ for $1 \leq i \leq m$ be the number of steps along the ith path. Then we have the following upper bound on the statistical dimension for the tangent cone $\mathcal{T}_{\mathcal{F}_{n,\alpha},\nu}$:*

$$\delta(\mathcal{T}_{\mathcal{F}_{n,\alpha},\nu}) \leq \log(en) \sum_{i=1}^{m} k^{(i)}.$$

Proof. Let $B_1^{(i)}, \ldots, B_{k^{(i)}}^{(i)}$ denote the contiguous blocks where $\nu^{(i)}$ is constant. For any vector $\theta \in \mathbb{R}^n$ and any $A \subset \{1 \leq i \leq n\}$ denote by θ_A the vector θ with coordinates restricted to be in the set A. Also denote the cone of nonincreasing sequences in $\mathbb{R}^d$ by $\mathcal{M}_d$. Equipped with this notation we now claim that

$$\mathcal{T}_{\mathcal{F}_{n,\alpha},\nu} \subset \mathcal{K} = \left\{\theta \in \mathbb{R}^n : \theta_{B_j^{(i)}} \in \mathcal{M}_{|B_j^{(i)}|} \text{ for all } 1 \leq i \leq m, 1 \leq j \leq k^{(i)}\right\}. \tag{3.8}$$

Assuming this claim for now, by monotonicity of statistical dimension, we have

$$\delta(\mathcal{T}_{\mathcal{F}_{n,\alpha},\nu}) \leq \delta(\mathcal{K}).$$

Since $\mathcal{K}$ is a cone composed of disjoint monotone pieces and the statistical dimension of the monotone cone is known to be $\delta(\mathcal{M}_d) \leq \log(ed)$ we have

$$\delta(\mathcal{T}_{\mathcal{F}_{n,\alpha},\nu}) \leq \delta(\mathcal{K}) \leq \sum_{i=1}^{m} \sum_{j=1}^{k^{(i)}} \log(e|B_j^{(i)}|) \leq \sum_{i=1}^{m} k^{(i)} \log(en).$$

It remains to prove (3.8). Take any element $\theta \in \mathcal{T}_{\mathcal{F}_{n,\alpha},\nu}$. Then by (3.7) there exists $t > 0$ and $v \in \mathcal{F}_{n,\alpha}$ such that $\theta = v - t\nu$. Now consider any block $B_j^{(i)}$. By definition, $\nu_{B_j^{(i)}}$ is a constant vector and $v_{B_j^{(i)}} \in \mathcal{M}_{|B_j^{(i)}|}$. This implies $\theta_{B_j^{(i)}} \in \mathcal{M}_{|B_j^{(i)}|}$, which proves (3.8). □

We are now ready to prove Theorem 2.4.

Proof of Theorem 2.4. Fix an arbitrary $\epsilon > 0$. Recall $l = \lceil n^{1-\alpha} \rceil$ and $m = \lceil n^{\alpha} \rceil$. For each path $\mu^{(i)} = (\mu_1^{(i)}, \dots, \mu_l^{(i)})$ we can use Lemma 3.6 to obtain a nonincreasing sequence $\nu^{(i)} = (\nu_1^{(i)}, \dots, \nu_l^{(i)})$ such that

$$k(\nu^{(i)}) \leq \frac{\mu_1^{(i)}}{\epsilon} + 1 \tag{3.9}$$

along with $\|\mu^{(i)} - \nu^{(i)}\|^2 \leq l\epsilon^2$ and $\nu_1^{(i)} \leq \mu_1^{(i)}$. Also define $\nu_1 = \mu_1$. Then it is clear that $\nu \in \mathcal{F}_{n,\alpha}$. Using Lemma (3.7) we deduce

$$\mathbb{E}\|\widehat{\mu} - \mu\|^2 \leq \|\nu - \mu\|^2 + \sigma^2 \delta(\mathcal{T}_{\mathcal{F}_{n,\alpha},\nu}).$$

Now we have $\|\nu - \mu\|^2 = \sum_{i=1}^m \|\nu^{(i)} - \mu^{(i)}\|^2 \leq ml\epsilon^2 = n\epsilon^2$. Therefore using (3.9) and Lemma 3.8 we can conclude that

$$\begin{aligned}
\mathbb{E}\|\widehat{\mu} - \mu\|^2 &\leq n\epsilon^2 + \sigma^2 \sum_{i=1}^m \Big(\frac{\mu_1^{(i)}}{\epsilon} + 1\Big) \log(en) \\
&\quad n\epsilon^2 + \sigma^2 \sum_{i=1}^m \Big(\frac{\mu_1^{(i)}}{\epsilon}\Big) \log(en) + \sigma^2 m \log(en) \\
&\leq n\epsilon^2 + \sigma^2 \frac{\mu_1}{\epsilon} \log(en) + \sigma^2 m \log(en).
\end{aligned}$$

By choosing $\epsilon = \big(\frac{\sigma^2 \mu_1 \log(en)}{n}\big)^{1/3}$ we obtain the risk bound

$$\mathbb{E}\|\widehat{\mu} - \mu\|^2 \leq 2n^{1/3}\big(\sigma^2 \mu_1 \log(en)\big)^{2/3} + \sigma^2 m \log(en),$$

which finishes the proof of the theorem. □

3.5. *Proof of Theorem 2.5: Simplex upper bound on the LSE for deep trees*

We begin by establishing some notation. If (M, ρ) is a metric space, a set $\widehat{T} \subset M$ is called an ϵ-cover of $T \subset M$ in case

$$T \subset \cup_{t \in \widehat{T}} B_\epsilon(t) \tag{3.10}$$

where $B_\epsilon(t) = \{x \in M : \rho(x, t) < \epsilon\}$ denotes the ball of radius $\epsilon > 0$ centered at $t \in M$. The ϵ-covering number of T is the cardinality of the smallest cover of T:

$$N(\epsilon, T, \rho) = \min\left\{|\widehat{T}| : \widehat{T} \text{ is an } \epsilon\text{-cover of } T\right\}. \tag{3.11}$$

The metric will almost always be given by the usual Euclidean norm, in which case we denote the covering numbers by $N(\epsilon, T)$; otherwise, the metric will be explicitly mentioned.

An important result for us here again is Theorem 3.1. It This now requires a bound on the expected supremum of the Gaussian process in the definition of $f_\mu(t)$ in (3.1). The following chaining result gives an upper bound on the expected supremum of the required Gaussian process. This chaining result is sometimes known as Dudley's entropy integral inequality; a proof for the version of the bound stated below can be found in Chatterjee et al. (2015a, Lemma A.2).

Theorem 3.9 (Chaining)**.** *For every $\mu \in \mathcal{F}(T)$ and $t > 0$,*

$$\mathbb{E}\left(\sup_{\nu \in \mathcal{F}(T): \|\nu - \mu\| \le t} \langle \nu - \mu, Z \rangle\right) \le \sigma \inf_{0 < \delta \le 2t} \left\{ 12 \int_\delta^{2t} \sqrt{\log N(\epsilon, B(\mu, t))}\, d\epsilon + 4\delta\sqrt{n} \right\}.$$

The first step in the proof of Theorem 2.5 is to upper bound the covering number for the metric space of flows.

Lemma 3.10. *Fix $V > 0$ and a positive integer n. For any tree T with n nodes and depth h denote by $\mathcal{F}_V(T) =: \mathcal{F}_V$ the set of flows on T where the root flow μ_1 is no greater than V. For any $\epsilon > 0$, define*

$$m_\epsilon = \left\lceil \frac{V^2 h}{\epsilon^2} \right\rceil. \tag{3.12}$$

Then

$$\log N(\epsilon, \mathcal{F}_V) \le m_\epsilon \left(1 + \log\left(1 + \frac{n}{m_\epsilon}\right)\right). \tag{3.13}$$

Proof. Let h denote the height of the tree, let $\mathcal{L} = \{i : \mathcal{C}(i) = \emptyset\} \subset [n]$ denote the set of leaf nodes, and let $\mathcal{I} = [n] - \mathcal{L}$ be the set of internal, or non-leaf nodes. We first note that a flow $\mu \in \mathcal{F}_V$ can be uniquely identified by the collection of *leaks* at the nodes, with the leak at node i defined by

$$l_i^\mu = \mu_i - \sum_{j \in \mathcal{C}(i)} \mu_j. \tag{3.14}$$

If $i \in \mathcal{L}$ is a leaf, then $l_i^\mu = \mu_i$ and we also call this the *residue* of the flow at i. By definition,

$$\sum_{i \in \mathcal{I} \cup \mathcal{L}} l_i^\mu \le \mu_1; \tag{3.15}$$

that is, the residues and leaks can together be no greater than the flow into the root node.

For any positive integer $m \in \mathbb{Z}_+$, we define a set of flows $\mathcal{F}_V^m \subset \mathcal{F}_V$ as

$$\mathcal{F}_V^m = \left\{ \mu \in \mathcal{F}_V : \frac{m}{V} l^\mu = (i_1, \ldots, i_n), \text{ for some } (i_1, \ldots, i_n) \in \mathbb{Z}_+^n \text{ with } i_1 + \cdots + i_n \le m \right\}. \tag{3.16}$$

We now show using a probabilistic argument that $\mathcal{F}_V^m$ is a $(V^2 h)/m$ covering set for $\mathcal{F}_V$. Fixing a flow $\mu \in \mathcal{F}_V$, we define a random flow F, whose distribution depends on μ, by specifying its

leaks l^F as follows:

$$l^F = \begin{cases} Ve_i & \text{with probability } \dfrac{l_i^\mu}{V} \\ 0 & \text{with probability } 1 - \dfrac{1}{V}\sum_i l_i^\mu. \end{cases} \tag{3.17}$$

Here e_i denotes the n dimensional vector with 1 in the ith entry and 0 everywhere else. For any positive integer m, we now let $F_1, \ldots, F_m$ be m i.i.d. copies of F, and define the mean flow

$$\overline{\mu}^m = \frac{1}{m}\left(F_1 + \cdots + F_m\right). \tag{3.18}$$

Note that

$$\overline{\mu}^m \in \mathcal{F}_V^m \text{ with probability } 1. \tag{3.19}$$

Consider now the expected Euclidean distance between μ and the random flow $\overline{\mu}^m$. We can write

$$\mathbb{E}\|\overline{\mu}^m - \mu\|^2 = \sum_{i=1}^n \mathbb{E}(\overline{\mu}_i^m - \mu_i)^2 \tag{3.20}$$

$$= \sum_i \text{Var}(\overline{\mu}_i^m), \tag{3.21}$$

where Var refers to the variance of a random variable. The last equality holds because $\mathbb{E}\overline{\mu}^m = \mu$; that is, $\overline{\mu}^m$ is unbiased for μ.

Let $\overline{l}_v^m$ denote the leak of the flow $\overline{\mu}^m$ at node v. Fix a leaf node $v \in \mathcal{L}$. Then the leak $\overline{l}_v^m$ is a mean of m Bernoulli random variables, each taking the value V with probability $\frac{\mu_v}{V}$ and 0 with the complementary probability. Hence we have that

$$\text{Var}(\overline{\mu}_v^m) = \frac{V^2}{m}\frac{\mu_v}{V}\left(1 - \frac{\mu_v}{V}\right) \leq \frac{V\mu_v}{m}. \tag{3.22}$$

More generally, we have that for any node v,

$$\overline{\mu}_v^m = \sum_{k \in \mathcal{C}(v)} \overline{\mu}_k^m + \overline{l}_v^m \tag{3.23}$$

and hence

$$\text{Var}(\overline{\mu}_v^m) = \text{Var}\left(\sum_{k \in \mathcal{C}(v)} \overline{\mu}_k^m + \overline{l}_v^m\right) \leq \sum_{k \in \mathcal{C}(v)} \text{Var}\left(\overline{\mu}_k^m\right) + \text{Var}(\overline{l}_v^m), \tag{3.24}$$

where the inequality holds since the random variables $\left(\overline{\mu}_1^m, \ldots, \overline{\mu}_1^m, \overline{l}_v^m\right)$ are pairwise negatively correlated, by construction of the random flow $\overline{\mu}^m$. Applying this argument recursively, we have

$$\text{Var}(\overline{\mu}_v^m) \leq \sum_{u \in \mathit{Sub}(v)} \text{Var}(\overline{l}_u^m), \tag{3.25}$$

where $Sub(v)$ denotes the subtree rooted at v. For any non-leaf $i \in \mathcal{I}$, by similar reasoning as in (3.22), we have

$$\text{Var}(\bar{l}_i^m) \leq \frac{V}{m} l_i^\mu. \tag{3.26}$$

Using this observation together with (3.21) and (3.25), and denoting by $d_i \leq h$ the depth of node i, we conclude

$$\begin{aligned}
\mathbb{E}\|\overline{\mu}^m - \mu\|^2 &= \sum_i \text{Var}(\overline{\mu}_i^m) && (3.27)\\
&\leq \sum_i d_i \text{Var}(\bar{l}_i^m) && (3.28)\\
&\leq \frac{V}{m} \sum_i d_i l_i^\mu && (3.29)\\
&\leq \frac{Vh}{m} \sum_i l_i^\mu && (3.30)\\
&\leq \frac{V^2 h}{m}. && (3.31)
\end{aligned}$$

The above result holds in expectation with respect to the random draw $\overline{\mu}^m$, which by (3.19) always lies in the finite set $\mathcal{F}_V^m$. Hence, we can assert the existence of an element $\widetilde{\mu} \in \mathcal{F}_V^m$ such that

$$\|\mu - \widetilde{\mu}\|^2 \leq \frac{V^2 h}{m}. \tag{3.32}$$

Since $\mu \in \mathcal{F}$ was arbitrarily chosen, we have shown that $\mathcal{F}_V^m$ is a $\left(V^2 h\right)/m$ covering set for $\mathcal{F}_V$.

Finally, note that the set $\mathcal{F}_V^m$ is in one-to-one correspondence with the set $\{(i_1, \ldots, i_n) \in \mathbb{Z}_+^n : i_1 + \cdots + i_n \leq m\}$. By standard combinatorics, we have

$$|\mathcal{F}_V^m| = \binom{n+m}{m} \leq \left(\frac{e(n+m)}{m}\right)^m. \tag{3.33}$$

For any given $\epsilon > 0$ we then choose $m = m_\epsilon$ according to (3.12) to deduce the statement of the lemma. □

Remark 3.1. The idea of the above proof to demonstrate covering sets by randomization is not new, and is sometimes referred to as "Maurey's argument." Our proof is a generalization of the proof in van der Vaart and Wellner (1996, Lemma 2.6.11) to the setting of flows on rooted trees.

We are now ready to prove Theorem 2.5.

Proof of Theorem 2.5. Fix $\mu \in \mathcal{F}$. We use Theorem 3.1 by first upper bounding the function

$$f_\mu(t) := \mathbb{E}\left(\sup_{\nu \in \mathcal{F}: \|\nu - \mu\| \leq t} \langle Z, \nu - \mu \rangle\right) - \frac{t^2}{2} \tag{3.34}$$

where $Z \sim N(0, \sigma^2 I_{n\times n})$, using Dudley's entropy integral inequality. Note that the diameter of the set $\{\nu \in \mathcal{F} : \|\nu - \mu\| \le t\}$ is at most $2t$. Setting $\delta = 1/\sqrt{n}$ in Theorem 3.9 we obtain

$$\mathbb{E}\left(\sup_{\nu\in\mathcal{F}:\|\nu-\mu\|\le t} \langle Z, \nu - \mu\rangle\right) \le 4\sigma + 12\sigma \int_{\frac{1}{\sqrt{n}}}^{2t} \sqrt{\log N(\epsilon, \{\nu \in \mathcal{F} : \|\nu - \mu\| \le t\})}d\epsilon \tag{3.35}$$

$$\le 4\sigma + 12\sigma \int_{\frac{1}{\sqrt{n}}}^{2t} \sqrt{\log N(\epsilon, \{\nu \in \mathcal{F} : \nu_1 \le \mu_1 + t\})}d\epsilon \tag{3.36}$$

where the second inequality follows from the inclusion

$$\{\nu \in \mathcal{F} : \|\nu - \mu\| \le t\} \subset \{\nu \in \mathcal{F} : \nu_1 \le \mu_1 + t\}. \tag{3.37}$$

Now we are in a position to use Lemma 3.10 since it gives us upper bounds on log covering numbers of sets of the form $\{\nu \in \mathcal{F} : \nu_1 \le \mu_1 + t\}$. In particular, we have

$$\log N(\epsilon, \{\nu \in \mathcal{F} : \nu_1 \le \mu_1 + t\})$$

$$\le \left(1 + \frac{(\mu_1 + t)^2 h}{\epsilon^2}\right)\left[1 + \log\left(1 + \frac{n}{1 + (\mu_1 + t)^2 h/\epsilon^2}\right)\right] \tag{3.38}$$

$$\le \frac{(\mu_1 + t)^2 h}{\epsilon^2}\left[1 + \log\left(1 + \frac{n\epsilon^2}{(\mu_1 + t)^2 h}\right)\right] + (1 + \log(n + 1)) \tag{3.39}$$

where we have used the facts that $\lceil x\rceil \le x+1$ and $x\,(1 + \log(1 + c/x))$ is a nondecreasing function of $x \in \mathbb{R}_+$ for $c > 0$. The simple inequality $\sqrt{a + b} \le \sqrt{a} + \sqrt{b}$ now gives

$$\mathbb{E}\left(\sup_{\nu\in\mathcal{F}:\nu_1\le\mu_1+t} \langle Z, \nu - \mu\rangle\right) \le 4\sigma + 12\sigma\left((\mu + t)\sqrt{h}I + 2t\sqrt{1 + \log(n + 1)}\right) \tag{3.40}$$

where I is the integral

$$I = \int_0^{2t} \frac{1}{\epsilon}\left(1 + \log\left[1 + \frac{n\epsilon^2}{(\mu_1 + t)^2 h}\right]\right)^{1/2} d\epsilon$$

$$\le C\left(1 + \log\left[1 + 4\frac{nt^2}{(\mu_1 + t)^2 h}\right]\right)^{3/2}$$

where the inequality follows from elementary calculus. We thus have

$$\mathbb{E}\Big(\sup_{\nu\in\mathcal{F}:\nu_1\le\mu_1+t} \langle Z, \nu - \mu\rangle\Big)$$

$$\le 4\sigma + 12\sigma C(\mu_1 + t)\sqrt{h}\left(1 + \log\left[1 + 4\frac{nt^2}{(\mu_1 + t)^2 h}\right]\right)^{3/2} + 24\sigma t\sqrt{1 + \log(n + 1)} \tag{3.41}$$

$$\le 4\sigma + 12\sigma C(\mu_1 + t)\sqrt{h}\left(1 + \log\left[1 + 4n\right]\right)^{3/2} + 24\sigma t\sqrt{1 + \log(n + 1)} \tag{3.42}$$

where the second inequality uses $t^2/(\mu_1 + t)^2 \le 1$ and $h \ge 1$. Therefore,

$$f_\mu(t) \le 4\sigma + 12\sigma C(\mu_1 + t)\sqrt{h}\left(1 + \log\left[1 + 4n\right]\right)^{3/2} + 24\sigma t\sqrt{1 + \log(n + 1)} - \frac{t^2}{2} \equiv g(t). \tag{3.43}$$

Letting t^* be the larger root of the quadratic function $g(t)$, we have

$$f_\mu(t^*) \leq g(t^*) = 0. \tag{3.44}$$

After some algebraic manipulation one can upper bound t^{*2} according to

$$t^{*2} \leq C\sigma^2 h(1+\log n)^3 + C\sigma\mu_1\sqrt{h}(1+\log n)^{3/2}.$$

Therefore equation (3.2) of Theorem 3.1 implies that

$$\mathbb{E}\|\widehat{\mu}-\mu\|^2 \leq C\left(\sigma^2 h(1+\log n)^3 + \sigma\mu_1\sqrt{h}(1+\log n)^{3/2}\right), \tag{3.45}$$

which completes the proof of the theorem. □

3.6. Proof of Theorem 2.7: Minimax lower bound for monotone sequences

We shall use Assouad's lemma to prove Theorem 2.7. The following version of Assouad's Lemma is a consequence of Lemma 24.3 of van der Vaart (2000).

Lemma 3.11 (Assouad). *Fix $V > 0$ and a positive integer d. Suppose that, for each $\tau \in \{-1,1\}^d$, there is an associated $g^\tau \in \mathcal{M}_V$. Then*

$$\inf_{\widetilde{\mu}} \sup_{\mu\in\mathcal{M}_V} \mathbb{E}\|\widetilde{\mu}-\mu\|^2 \geq \frac{d}{8}\min_{\tau\neq\tau'}\frac{\|g^\tau - g^{\tau'}\|^2}{\Upsilon(\tau,\tau')}\min_{\Upsilon(\tau,\tau')=1}\left(1-\|\mathbb{P}_{g^\tau}-\mathbb{P}_{g^{\tau'}}\|_{TV}\right),$$

where $\Upsilon(\tau,\tau') := \sum_{i=1}^d I\{\tau_i \neq \tau_i'\}$ denotes the Hamming distance between τ and τ' and $\|\cdot\|_{TV}$ denotes the total variation distance. The notation $\mathbb{P}_g$ for $g \in \mathcal{M}_V$ refers to the joint distribution of $y_i = g_i + \epsilon_i$, for $1 \leq i \leq n$ when ϵ_i are independent normally distributed random variables with mean zero and variance σ^2.

Proof of Theorem 2.7. Fix any integer $1 \leq k \leq n$ and define $m = \lfloor n/k \rfloor$. Define the vector $\mu_i^* = Vi/n$. For any $\tau \in \{0,1\}^k$ define the vector μ^τ in the following manner:

$$\mu_j^\tau = \begin{cases} \mu^*_{(i-1)m+1} & \text{if } \tau_i = 0, (i-1)m < j \leq im \\ \mu^*_{im} & \text{if } \tau_i = 1, (i-1)m < j \leq im. \end{cases} \tag{3.46}$$

If $j > mk$, then define $\mu_j^\tau = V$ for any $\tau \in \{0,1\}^k$. It is clear that $\mu\tau \in \mathcal{M}_V$. Now note

$$\|\mu^\tau - \mu^{\tau'}\|^2 = \frac{V^2}{n^2}\Upsilon(\tau,\tau')(1^2+2^2+\cdots+m^2).$$

It is now not hard to see that

$$\frac{V^2m^3}{3n^2}\Upsilon(\tau,\tau') \leq \|\mu^\tau - \mu^{\tau'}\|^2 = \frac{V^2m^3}{n^2}\Upsilon(\tau,\tau'). \tag{3.47}$$

Also by Pinsker's inequality we have

$$(\|\mathbb{P}_{\mu^\tau} - \mathbb{P}_{\mu^{\tau'}}\|_{TV})^2 \leq \frac{1}{2}\boldsymbol{KL}(\mathbb{P}_{\mu^\tau}, \mathbb{P}_{\mu^{\tau'}}) = \frac{1}{4\sigma^2}\|\mu^\tau - \mu^{\tau'}\|^2 \leq \frac{V^2m^3}{4\sigma^2n^2}\Upsilon(\tau, \tau'),$$

where we used (3.47) in the last inequality. An application of Assouad's Lemma, along with the last two equations, now yields the minimax lower bound

$$\inf_{\widetilde{\mu}} \sup_{\mu\in\mathcal{M}_V} \mathbb{E}\|\widetilde{\mu} - \mu\|^2 \geq \frac{k}{8}\frac{V^2m^3}{3n^2}\Big(1 - \sqrt{\frac{V^2m^3}{4\sigma^2n^2}}\Big) \geq \frac{V^2m^2}{48n}\Big(1 - \sqrt{\frac{V^2m^3}{4\sigma^2n^2}}\Big) = r_m,$$

where we have used $km \geq n/2$ in the last inequality.

Note that the last equation gives a minimax lower bound depending on m, which can be chosen to be any positive integer not bigger than n. Observe, however, that choosing $m = 1$ results in the degenerate case where all the μ^τ defined in (3.46) are the same vector. We can therefore write the minimax lower bound in the form

$$R_{n,V,\sigma} = \inf_{\widetilde{\mu}} \sup_{\mu\in\mathcal{M}_V} \mathbb{E}\|\widetilde{\mu} - \mu\|^2 \geq \max_{2\leq m\leq n} \frac{V^2m^2}{48n}\Big(1 - \sqrt{\frac{V^2m^3}{4\sigma^2n^2}}\Big). \tag{3.48}$$

We now have three cases to investigate.

1. $nV^2 \leq \min\{n\sigma^2, (\sigma^2)^{2/3}(nV^2)^{1/3}\}$: In this case, we have $nV^2 \leq \sigma^2$. In this case set $m = n$ in the right side of (3.48) to get the minimax lower bound

$$R_{n,V,\sigma} \geq \frac{nV^2}{96}.$$

2. $n\sigma^2 \leq \min\{nV^2, (\sigma^2)^{2/3}(nV^2)^{1/3}\}$: In this case we have $n\sigma \leq V$. This implies $R_{n,V,\sigma} \geq R_{n,n\sigma/(2\sqrt{2}),\sigma}$. Now a lower bound for $R_{n,n\sigma/(2\sqrt{2}),\sigma}$ can be obtained by setting $m = 2$ in (3.47) to obtain

$$R_{n,V,\sigma} \geq \frac{n\sigma^2}{192}.$$

3. $(\sigma^2)^{2/3}(nV^2)^{1/3} < \min\{n\sigma^2, nV^2\}$: In this case we have $V < n\sigma$ and $\sigma^2 < nV^2$. We further subdivide this case into two subcases.

   a) Suppose $\sigma^2n^2/V^2 < 8$. Then we can write

$$R_{n,V,\sigma} \geq R_{n,n\sigma/(2\sqrt{2}),\sigma} \geq \frac{n\sigma^2}{192} \geq \frac{(\sigma^2)^{2/3}(nV^2)^{1/3}}{192}.$$

   The first inequality follows from setting $m = 2$ in (3.47) and the second inequality follows because $(\sigma^2)^{2/3}(nV^2)^{1/3} \leq n\sigma^2$.

   b) Suppose $\sigma^2n^2/V^2 \geq 8$. We now set $m = \lfloor\big(\sigma^2n^2)/V^2\big)^{1/3}\rfloor$. We then have $m \geq 2$ and also we must have $m \leq n$ because $\sigma^2 < nV^2$. Hence, m is a feasible choice to be used in (3.47). We now use our choice of m in (3.47) to obtain

$$R_{n,V,\sigma} \geq \frac{(\sigma^2)^{2/3}(nV^2)^{1/3}}{192}.$$

Here we have used the fact that $m \geq \left(\sigma^2 n^2)/V^2\right)^{1/3}/2$ and $\frac{V^2 m^3}{4\sigma^2 n^2} \geq 1/4$. The first fact is true because $m \geq 2$ and the second fact follows trivially from the definition of m.

This finishes the proof of the theorem. □

3.7. *Proof of Theorem 2.8: Minimax lower bound for deep trees,* $\alpha \leq \frac{1}{3}$

Proof of Theorem 2.8. Denote the subset of flows in $\mathcal{F}(\mathcal{T}_{\alpha,n})$ with root value at most V by $\mathcal{F}_{n,\alpha,V}$. Recall $m = \lceil n^\alpha \rceil$. Fix any $(v_1, \ldots, v_m)$ such that $\sum_{i=1}^m v_i \leq V$. Consider the space of flows $\mathcal{F}(n, V, (v_1, \ldots, v_m))$ where the root is set at V and its m children are set at $(v_1, \ldots, v_m)$. Clearly $\mathcal{F}(n, V, (v_1, \ldots, v_m)) \subset \mathcal{F}_{n,\alpha,V}$. Hence we have

$$\inf_{\widetilde{\mu}} \sup_{\mu \in \mathcal{F}_{n,\alpha,V}} \mathbb{E}\|\widetilde{\mu} - \mu\|^2 \geq \inf_{\widetilde{\mu}} \sup_{\mu \in \mathcal{F}(n,V,(v_1,\ldots,v_m))} \mathbb{E}\|\widetilde{\mu} - \mu\|^2,$$

and it suffices to lower bound the right side of the above equation. Now the estimation problem in $\mathcal{F}(n, V, (v_1, \ldots, v_m))$ is just m separate isotonic regression problems. This means that the inf sup term over $\mathcal{F}(n, V, (v_1, \ldots, v_m))$ decomposes into a sum of inf sup terms over each of the paths with monotonicity constraints. Applying the minimax lower bound for isotonic regression given in Theorem 2.7 to each of these subproblems completes the proof. □

Proof of Corollary 2.9. Set $v_i = \sigma V/n^\alpha$ in the lower bound given by Theorem 2.8. Since $n^{3\alpha-1} \leq V \leq n$,

$$\min\{n^{1-\alpha}\sigma^2, (v_i\sigma^2)^{2/3} n^{(1-\alpha)/3}, n^{1-\alpha} v_i^2\} = (v_i\sigma^2)^{2/3} n^{(1-\alpha)/3}.$$

Theorem 2.8 then implies that a valid minimax lower bound in terms of the sum of squared errors $n^\alpha (v_i\sigma^2)^{2/3} n^{(1-\alpha)/3} = (V\sigma^2)^{2/3} n^{1/3}$ by the choice of v_i. □

3.8. *Proof of Theorem 2.10: Minimax lower bound for deep trees,* $\alpha \geq \frac{1}{3}$

The proof of Theorem 2.10 is very similar to that of Theorem 2.2, where we use the Varshamov-Gilbert Lemma 3.3 and Fano's lemma 3.2. The details are omitted.

3.9. *Proof of Theorem 2.11: Tightness of the minimax lower bound,* $\alpha \geq \frac{1}{3}$

The following lemma gives an upper bound to the minimax rate in a general Gaussian denoising problem where the mean is known to lie in a set K. The upper bound is information-theoretic and is expressed in terms of covering numbers of K.

Theorem 3.12. *Let $y \sim N(\theta^*, \sigma^2 I)$ be a n dimensional random vector. Let $K \subset \mathbb{R}^n$ and $\theta^* \in K$. Then*

$$\inf_{\widetilde{\theta}} \sup_{\theta^* \in K} \mathbb{E}\|\widetilde{\theta} - \theta^*\|^2 \leq \inf_{\epsilon > 0} \left(16\sigma^2 \log N(\epsilon, K, \|.\|) + 3\epsilon^2\right).$$

Proof. Let $F \subset K$ be a finite subset. Define the least squares estimator over the finite set F as

$$\widehat{\theta}_F = \operatorname*{argmin}_{\mu \in F} \|y - \mu\|^2.$$

We start with the following inequality which holds for any nonnegative function $G : F \to \mathbb{R}_+$:

$$G(\widehat{\theta}_F) \leq \sum_{\mu \in \mathcal{F}} G(\mu) \exp\left(\alpha \|y - \widehat{\theta}_F\|^2 - \alpha \|y - \mu\|^2\right).$$

Because $\widehat{\theta}_F$ is the least squares estimator, we can replace it by any arbitrary but fixed and data-independent $\theta' \in F$ in the right side of the above inequality. Then taking expectations on both sides we obtain

$$\mathbb{E}\, G(\widehat{\theta}_F) \leq \sum_{\mu \in \mathcal{F}} G(\mu) \mathbb{E} \exp\left(\alpha \|y - \theta'\|^2 - \alpha \|y - \mu\|^2\right). \tag{3.49}$$

Writing $y = \theta^* + z$ where $z \sim N(0, \sigma^2 I)$, some elementary algebra gives us

$$\|y - \theta'\|^2 - \|y - \mu\|^2 = \|\theta^* - \theta'\|^2 - \|\theta^* - \mu\|^2 + 2\langle z, \mu - \theta' \rangle.$$

Knowing the moment generating function of z then lets us conclude

$$\mathbb{E} \exp\left(\alpha \|y - \theta'\|^2 - \alpha \|y - \mu\|^2\right) = \exp\left(\alpha \|\theta^* - \theta'\|^2 - \alpha \|\theta^* - \mu\|^2 + 2\alpha^2 \sigma^2 \|\mu - \theta'\|^2\right).$$

The elementary inequality $\|\mu - \theta'\|^2 \leq 2\|\mu - \theta^*\|^2 + 2\|\theta^* - \theta'\|^2$ applied to the last equation, together with (3.49), gives us

$$\mathbb{E}\, G(\widehat{\theta}_F) \leq \sum_{\mu \in \mathcal{F}} G(\mu) \exp\left((4\alpha^2\sigma^2 - \alpha)\|\mu - \theta^*\|^2 + (4\alpha^2\sigma^2 + \alpha)\|\theta^* - \theta'\|^2\right). \tag{3.50}$$

The choices $G(\mu) = \exp(\frac{\|\mu - \theta^*\|^2}{16\sigma^2})$ and $\alpha = \frac{1}{8\sigma^2}$ then establish the risk bound

$$\mathbb{E} \exp(\frac{\|\widehat{\theta}_F - \theta^*\|^2}{16\sigma^2}) \leq |F| \exp(\frac{3}{16\sigma^2}\|\theta^* - \theta'\|^2).$$

Using Jensen's inequality on the left side and taking logarithms yields

$$\mathbb{E}\, \frac{\|\widehat{\theta}_F - \theta^*\|^2}{16\sigma^2} \leq \log|F| + \frac{3}{16\sigma^2}\|\theta^* - \theta'\|^2.$$

Since $\theta' \in F$ was arbitrary we can actually conclude

$$\mathbb{E}\, \frac{\|\widehat{\theta}_F - \theta^*\|^2}{16\sigma^2} \leq \log|F| + \min_{\mu \in F} \frac{3}{16\sigma^2}\|\theta^* - \mu\|^2.$$

Now, if $F = F_\epsilon$ is chosen to be a ϵ cover for K, we have

$$\mathbb{E}\, \frac{\|\widehat{\theta}_F - \theta^*\|^2}{16\sigma^2} \leq \log N(\epsilon, K, \|.\|) + \frac{3}{16\sigma^2}\epsilon^2.$$

Taking the infimum over $\epsilon > 0$ finishes the proof. □

Remark 3.2. *This basic idea of the above result can be traced back to the paper Barron et al. (2008) and the references therein. A more general version of the above theorem can be found in Chatterjee (2014a) (Theorem 1.2.2).*

Proof of Theorem 2.11. We use Lemma 3.10 to get an upper bound on the log covering number of $\mathcal{F}_{n,\alpha,V}$ of the form

$$\log N(\epsilon, \mathcal{F}_{n,\alpha,V}, \|.\|) \leq (1 + \frac{V^2 n^{1-\alpha}}{\epsilon^2})(1 + \log(en)).$$

An application of Theorem 2.11 immediately yields an upper bound to the minimax rate R as

$$\frac{R}{16\sigma^2} \leq \inf_{\epsilon>0} \big((1 + \frac{V^2 n^{1-\alpha}}{\epsilon^2})(1 + \log(en)) + \frac{3}{16\sigma^2}\epsilon^2\big).$$

Setting $\epsilon = \big(\frac{16\sigma^2 V^2 n^{1-\alpha}\log(en)}{3}\big)^{1/4}$ we obtain

$$R \leq 12\big(\frac{16\sigma^2 n^{1-\alpha}\log(en)}{3}\big)^{1/2} \leq 21\, V\, \sigma\, n^{(1-\alpha)/2}\log(en)^{1/2} + 16\sigma^2\big(1 + \log(en)\big)$$

as an upper bound to the risk. □

3.10. Proof of Theorem 2.12: Tightness of the LSE upper bound, $\alpha \geq \frac{1}{3}$

Proof. We will again prove the lower bound to the risk at the origin, taking $\mu = 0$. Lemma 3.5 implies the pointwise inequality

$$\|\widehat{\mu} - \mu\| = \sup_{\nu \in \mathcal{F}_{n,\alpha}: \|\nu\| \leq 1} \langle Z, \nu \rangle.$$

Therefore we can now write

$$\big(\mathbb{E} \sup_{\nu \in \mathcal{F}_{n,\alpha}: \|\nu\| \leq 1} \langle Z, \nu \rangle\big)^2 \leq \mathbb{E}\big(\sup_{\nu \in \mathcal{F}_{n,\alpha}: \|\nu\| \leq 1} \langle Z, \nu \rangle\big)^2 = \mathbb{E}\|\widehat{\mu} - \mu\|^2. \tag{3.51}$$

Thus, it suffices to lower bound the Gaussian width term $\mathbb{E}\sup_{\nu \in \mathcal{F}_{n,\alpha}: \|\nu\| \leq 1} \langle Z, \nu \rangle$.

Consider the case $1/3 \leq \alpha \leq 1/2$. For $1 \leq i \leq m$ define $S^{(i)} = \sum_{j=1}^{l} Z_j^{(i)}$ and define the random signs

$$s^{(i)} = sign(S^{(i)}).$$

Now define a random flow v in the following fashion. For each $1 \leq i \leq m$ and $1 \leq j \leq l$ define

$$v_j^{(i)} = s^{(i)} \frac{1}{2\sqrt{n}}.$$

Also define $v_1 = 1/2$. It is easy to check that $v \in \mathcal{F}_{n,\alpha}$ and $\|v\| \leq 1$. Therefore, we can write

$$\begin{aligned}\mathbb{E} \sup_{\nu \in \mathcal{F}_{n,\alpha}:\|\nu\|\leq 1} \langle Z, \nu \rangle \geq \mathbb{E}\langle Z, v\rangle &= \mathbb{E}\big(\frac{Z_1}{2} + \frac{1}{2\sqrt{n}} \sum_{i=1}^{m} |S^{(i)}|\big) \\ &= \frac{1}{2\sqrt{n}} \sum_{i=1}^{m} \mathbb{E}|S^{(i)}| \\ &= \frac{m}{2\sqrt{n}} 2\pi\sqrt{l}\sigma \\ &= \pi\sigma n^{\alpha/2}.\end{aligned}$$

This is because $S^{(i)} \sim N(0, l\sigma^2)$ and $\mathbb{E}|Z| = 2\pi\sigma$ where $Z \sim N(0, \sigma^2)$. Using (3.51) then allows us to conclude

$$\mathbb{E}\|\widehat{\mu} - \mu\|^2 \geq \pi^2\sigma^2 n^{\alpha}. \tag{3.52}$$

Now let us consider the case when $1/2 \leq \alpha \leq 1$. Define a random flow v as follows. For each $1 \leq i \leq m$ define

$$v_j^{(i)} = s^{(i)} \frac{1}{2\sqrt{n}} \mathcal{I}\{i \leq \sqrt{n}\}.$$

Also define $v_1 = 1/2$. It is again easy to check that $v \in \mathcal{F}_{n,\alpha}$ and $\|v\| \leq 1$. Hence we have

$$\begin{aligned}\mathbb{E} \sup_{\nu \in \mathcal{F}_{n,\alpha}:\|\nu\|\leq 1} \langle Z, \nu \rangle \geq \mathbb{E}\langle Z, v\rangle &= \mathbb{E}\big(\frac{Z_1}{2} + \frac{1}{2\sqrt{n}} \sum_{i=1}^{\sqrt{n}} |S^{(i)}|\big) \\ &= \frac{1}{2\sqrt{n}} \sum_{i=1}^{\sqrt{n}} \mathbb{E}|S^{(i)}| \\ &= \frac{\sqrt{n}}{2\sqrt{n}} 2\pi\sqrt{l}\sigma = \pi\sigma n^{(1-\alpha)/2}.\end{aligned}$$

Now using (3.51) lets us conclude that

$$\mathbb{E}\|\widehat{\mu} - \mu\|^2 \geq \pi^2\sigma^2 n^{(1-\alpha)},$$

completing the proof of the theorem. □

4. Simulations

In this section we present results from simulations to gain a qualitative understanding of the rates of convergence of the least squares estimator. We investigate the performance of the LSE in the family of trees $\mathcal{T}_{n,\alpha}$ for various values of α. Note that $\alpha = 1$ corresponds to the star graph. For each tree, we repeat the denoising experiment 100 times with sample size n growing from 9,000 to 20,000 in increments of 1000. In each experiment we obtain the squared error $\|\widehat{\mu} - \mu\|^2$; hence for each sample size our estimate of the expected sum of squares $\mathbb{E}\|\widehat{\mu} - \mu\|^2$ is an average of 100 trials. We take the log of these estimates and fit a linear regression. The slope is an estimate of the exponent of increase of $\mathbb{E}\|\widehat{\mu} - \mu\|^2$, as we expect $\mathbb{E}\|\widehat{\mu} - \mu\|^2$ to increase like n^{ν} for some ν.

We performed simulations for $\alpha = 0.4, 0.5, 0.75, 1$. The true flow μ was selected by setting the root vertex have value $\mu_1 = 1$. For any α, the children of the root were set to have value $\frac{1}{n^\alpha}$ and then the paths decreased from $\frac{1}{n^\alpha}$ to 0 in equal increments. Such a flow μ was chosen because in the case of isotonic regression ($\alpha = 0$), the mean vector which increases linearly from 0 to 1 in increments of $1/n$ has LSE with error scaling according to $n^{1/3}$, which is the worst case behavior; see Chatterjee et al. (2015b). Below we show plots (on a log scale) for the empirical squared error, averaged over 100 trials, versus sample size.

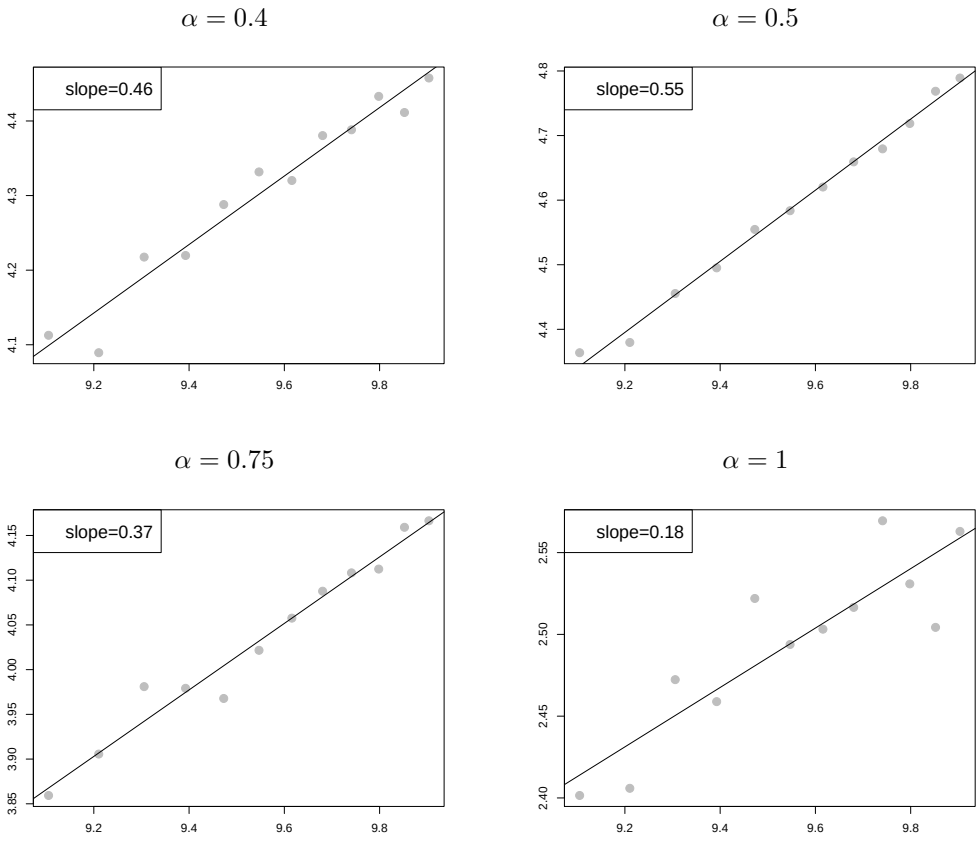

FIG 4. *For each value of α, a total of 100 denoising experiments were run with the size n of the tree growing from $9,000$ to $20{,}000$ in increments of $1{,}000$. In each experiment we obtain the squared error $\|\widehat{\mu} - \mu\|^2$, and average over the trials. The slope is an estimate of the exponent of increase of $\mathbb{E}\|\widehat{\mu} - \mu\|^2$, as the analysis shows that the expected squared error increases n^ν for some ν. The simulations are consistent with the theoretical findings that the risk bound behaves like $n^{1/2}$ around $\alpha = 0.5$, with the exponent in the bound decreasing linearly for α between 0.5 and 1.*

Note that the upper bound to the expected squared error in our theoretical analysis gives the exponent 0.4, 0.5, 0.25 for $\alpha = 0.4$, 0.5 and 0.75, respectively, and gives logarithmic growth for $\alpha = 1$. The estimated slopes for $\alpha = 0.4$ and 0.5 are close to 0.4 and 0.5, while the slopes are slightly higher than our analysis shows $\alpha = 0.75$ and 1. This may be because the logarithmic factors in our risk bound inflates the simulated slope.

We have simulated only for α larger than 0.4. This is because to be able to observe the correct rates in simulations for small α would require prohibitively large sample sizes. For example, with $\alpha = 0.1$ and $n < 20{,}000$ note that $n^{0.1} < 3$. Hence, the risk will behave very similarly to three separate isotonic regressions—we would observe rates of convergence scaling like $n^{(1-0.1)/3} = n^{0.3}$. In order to truly reflect the fact that the number of children of the root is growing like $n^{0.1}$ would require impractically large sample sizes. We observe, however, that a more computationally tractable estimator for small α is $\widehat{\mu}_{\text{nat}}$, which fits the top part of the tree, which is an $(m+1)$-node star, conditions on the estimated values of the children, and then fits m separate isotonic regressions. This algorithms scales linearly.

We have performed the simulations using a generic QP solver from Mosek, and the code takes slightly less than a minute for one denoising experiment with $n = 20{,}000$. However, the computation grows quickly for n much larger than $20{,}000$. As shown in the plots, the simulations are consistent with our theoretical findings.

5. Discussion

We have formulated a normal means problem on rooted trees, where the constraints on the means mimic a fluid flowing down from the root with possible leakage. We have studied the least squares estimator in this setting, showing risk bounds that only depend on the height of the tree. For trees with bounded or logarithmically growing diameter, this gives nearly sharp bounds, which match the minimax lower bounds. We have also studied the flow estimation problem for a family of trees that interpolates between a single path and a star graph. Here we find that the rate of convergence of the risk for the LSE is not monotonic in the depth of the tree. Moreover, we find a gap between the LSE and the minimax rates. To obtain matching upper and lower bounds over this family of trees for both the LSE and minimax estimators, we employ a range proof techniques, as summarized in Table 1. Our findings are displayed graphically in Figure 1.

A natural direction for future study is automatic adaptivity of the LSE to flows that are piecewise constant on long paths. Such adaptivity has recently been shown to hold for isotonic regression (Chatterjee et al., 2015b). Note that we have already obtained a result in this direction in Theorem 2.4, the proof of which proceeds by first showing fast rates for flows that are piecewise constant, and then using an approximation argument. It would also be interesting to pursue fast algorithms to denoise flows on rooted trees. In particular, one could explore connections to trend filtering (Wang et al., 2014), and the use of penalization and convex relaxations for flow estimation. For the path graph, the well known pooled adjacent violators algorithm gives a linear time procedure (Robertson et al., 1988). In case the tree is a star graph with root and n children, the problem of computing the LSE is related to computing the projection onto a probability simplex in $\mathbb{R}^n$. It may thus be possible to compute the LSE for certain trees $T_{\alpha,n}$ in $O(n)$ time by slightly modifying the algorithm in Duchi et al. (2008). But our results also present the challenge of closing the gap between the LSE and the minimax lower bound with computationally efficient estimators.

Acknowledgements

Research supported in part by ONR grant 11896509 and NSF grant DMS-1513594.

References

Amelunxen, D., Lotz, M., McCoy, M. B., and Tropp, J. A. (2013). Living on the edge: A geometric theory of phase transitions in convex optimization. *arXiv preprint arXiv:1303.6672*.

Ayer, M., Brunk, H. D., Ewing, G. M., Reid, W. T., and Silverman, E. (1955). An empirical distribution function for sampling with incomplete information. *Ann. Math. Statist.*, 26:641–647.

Barron, A. R., Huang, C., Li, J. Q., and Luo, X. (2008). Mdl, penalized likelihood, and statistical risk. In *Information Theory Workshop, 2008. ITW'08. IEEE*, pages 247–257. IEEE.

Bellec, P. C. (2015). Sharp oracle inequalities for least squares estimators in shape restricted regression. *arXiv preprint arXiv:1510.08029*.

Benabbas, S., Lee, H. C., Oren, J., and Ye, Y. (2011). Efficient sum-based hierarchical smoothing under\ ell_1-norm. *arXiv preprint arXiv:1108.1751*.

Birgé, L. and Massart, P. (1993). Rates of convergence for minimum contrast estimators. *Probab. Theory Related Fields*, 97(1-2):113–150.

Brunk, H. D. (1955). Maximum likelihood estimates of monotone parameters. *Ann. Math. Statist.*, 26:607–616.

Chatterjee, S. (2014a). *Adaptation in Estimation and Annealing*. PhD thesis, Yale University.

Chatterjee, S. (2014b). A new perspective on least squares under convex constraint. *The Annals of Statistics*, 42(6):2340–2381.

Chatterjee, S., Guntuboyina, A., and Sen, B. (2015a). On matrix estimation under monotonicity constraints. *arXiv preprint arXiv:1506.03430*.

Chatterjee, S., Guntuboyina, A., Sen, B., et al. (2015b). On risk bounds in isotonic and other shape restricted regression problems. *The Annals of Statistics*, 43(4):1774–1800.

Donoho, D. (1991). Gelfand n-widths and the method of least squares. Technical report, University of California, Berkeley. Department of Statistics.

Duchi, J., Shalev-Shwartz, S., Singer, Y., and Chandra, T. (2008). Efficient projections onto the l 1-ball for learning in high dimensions. In *Proceedings of the 25th international conference on Machine learning*, pages 272–279. ACM.

Graham, S. L., Kessler, P. B., and McKusick, M. K. (1982). gprof: A call graph execution profiler.

In *Proceedings of the ACM SIGPLAN 1982 Symposium on Compiler Construction*. SIGPLAN Notices 17, 6.

Kim, S.-J., Koh, K., Boyd, S., and Gorinevsky, D. (2009). ℓ_1 trend filtering. *SIAM review*, 51(2):339–360.

Meyer, M. and Woodroofe, M. (2000). On the degrees of freedom in shape-restricted regression. *Ann. Statist.*, 28(4):1083–1104.

Robertson, T., Wright, F. T., and Dykstra, R. L. (1988). *Order restricted statistical inference*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Ltd., Chichester.

Sharpnack, J. (2013). *Graph Structured Normal Means Inference*. PhD thesis, Carnegie Mellon University.

Skinner, C. J., Holt, D., and Smith, T. M. F. (1989). *Analysis of Complex Surveys*. Wiley Series in Probability and Mathematical Statistics. Wiley.

Spivey, J. M. (2004). Fast, accurate call graph profiling. *Software: Practice & Experience*, 34.

Tibshirani, R. J. et al. (2014). Adaptive piecewise polynomial estimation via trend filtering. *The Annals of Statistics*, 42(1):285–323.

van de Geer, S. (1990). Estimating a regression function. *Ann. Statist.*, 18(2):907–924.

van de Geer, S. (1993). Hellinger-consistency of certain nonparametric maximum likelihood estimators. *Ann. Statist.*, 21(1):14–44.

van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.

van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence and Empirical Process: With Applications to Statistics*. Springer-Verlag.

van Eeden, C. (1958). *Testing and estimating ordered parameters of probability distributions.* Mathematical Centre, Amsterdam.

Wang, Y. (1996). The l_2 risk of an isotonic estimate. *Comm. Statist. Theory Methods*, 25:281–294.

Wang, Y.-X., Sharpnack, J., Smola, A., and Tibshirani, R. J. (2014). Trend filtering on graphs. *arXiv preprint arXiv:1410.7690*.

Zhang, C.-H. (2002). Risk bounds in isotonic regression. *Ann. Statist.*, 30(2):528–555.